

Fair Allocation of Indivisible Public Goods

Brandon Fain*

Kamesh Munagala[†]

Nisarg Shah[‡]

Abstract

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 group-wise 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. More generally, if the feasibility constraints define 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 variants of the convex program for maximizing the Nash social welfare, but differ significantly from previous work in how it is used. Our guarantees are meaningful even when there are fewer elements than the number of agents. As far as we are aware, our work is the first to approximate the core in indivisible settings.

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 [25]. 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

*Department of Computer Science, Duke University, Durham, NC 27708. btfain@cs.duke.edu.

[†]Department of Computer Science, Duke University, Durham NC 27708. kamesh@cs.duke.edu.

[‡]Department of Computer Science, University of Toronto, Toronto ON, M5S3G8, Canada. nisarg@cs.toronto.edu

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?

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 [10], 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 [24, 10, 11, 3, 12, 9, 30]. 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* \mathbf{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 $\mathbf{c} \in \mathcal{F}$ is $u_i(\mathbf{c}) = \sum_{j \in \mathbf{c}} u_{ij}$. 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: $u_i(\mathbf{c})$ can be as large as m . Specifically, let $V_i = \max_{\mathbf{c} \in \mathcal{F}} u_i(\mathbf{c})$, 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 [22] for a formal introduction to matroids).

This generalizes the *public decision making* setting introduced by [10]. 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 \mathbf{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 in turn generalizes the classical setting of *private goods allocation* [24, 9, 30] in which private goods must be divided among agents with additive utilities, with each good 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 $A\mathbf{x} \leq \mathbf{b}$, 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 [31]:

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

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 [29, 18, 17, 11, 15, 12, 6], 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 that 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. *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 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, [10] 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 [24, 9]. 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 guarantee only when the number of issues is larger than the number

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.

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 [10]:

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

Below, we define the notion of *core outcomes*, which has been extensively studied (in similar forms) as a notion of stability in economics [13, 34, 26] and computer science [23, 11] in the context of *randomized* or *fractional* allocations. Our main contribution is to study it in the context of *integer* allocations.

Definition 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 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 1 (Section 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. *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. provide to individual agents [10].

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

Theorem 2. *If feasible outcomes are constrained to be matchings in an undirected graph, then for constant $\delta \in (0, 1]$, a $(\delta, 8 + 6/\delta)$ -core outcome can be computed in time polynomial in n and m .*

Our results in Section 5 are 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 and 2; we present an outline of the techniques in Section 5.2.

Theorem 3. *For constant $\delta \in (0, 1)$, given K packing constraints $\mathbf{Ax} \leq \mathbf{b}$ 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 — either could be large with the other two being small. 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 Example 2 (Appendix A), 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.

Finally, in Appendix C, 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. We show that 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 [34]. It has been extensively studied in public goods settings [13, 26, 11]. Most literature so far has considered the core with *fractional allocations*. Our definition of core (Definition 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 [27]. 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, [10] 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 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) and matching constraints (Theorem 2), 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 agent endowments, not scaled utilities. In more detail, in Definition 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 [34, 13] and has been considered in the context of participatory budgeting [11] 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 total 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 on the Lindahl equilibrium [13], 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 further, it is difficult to define such a notion of endowments in settings such as matroid or matching constraints. In the context of integer allocations with packing constraints, we extend our techniques to provide approximations to the notion of core with endowments in Appendix C, though this is not the main focus of our paper.

The notion of core with endowments logically implies a number of fairness notions considered in multi-winner election literature, such as justified representation, extended justified representation [3], and proportional justified representation [33]. 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 and 2 are similar to local search algorithms for proportional approval voting (PAV) [36, 4] that achieve proportional representation. It would be interesting to explore the connection between these various notions in greater depth.

Private Goods and Envy-freeness. Private goods are a special case of public goods with matroid constraints. Fair allocation of private goods is a widely studied topic [37, 16, 28, 21]. A common fairness criterion for private goods is envy-freeness: that no agent should (strongly) prefer the allocation of another agent. For fractional allocations, the classic context for envy-free allocation is cake cutting [32, 5]. 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 [24, 8, 9, 30], 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 [20].

Strategyproofness. In this work, we will 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, [35] shows that the only strategyproof and Pareto efficient mechanisms are dictatorial, and thus highly unfair, even when there

are only two agents with additive utilities over divisible goods. Second, our work is motivated by public goods settings with a large number of agents, such as participatory budgeting, wherein individual agents often have limited influence over the final outcome. It would be interesting to establish this formally, using 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 [9, 10] 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 .

[10] 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 1. Consider an instance of public decision making [10] 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 $\mathbf{c} = (a_1^1, a_1^2, \dots, a_1^m)$ 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 [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). In contrast, in Example 3 (Appendix A) we show 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:

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

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 Appendix B.1.

Theorem 4. 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.

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 $F(\mathbf{c}') - F(\mathbf{c}) \geq \frac{n\gamma}{m}$.
 - 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 1 ([22]). *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 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 goods 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 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.

Lemma 3. For private goods allocation, maximizing the smooth Nash welfare objective with $\ell = 1$ returns a $(0, 1)$ -core outcome.

Proof. There is a set of agents N and a set of private goods M . Each agent $i \in N$ has a utility function $u_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$. Utilities are additive, so $u_i(S) = \sum_{g \in S} u_i(\{g\})$ for all $S \subseteq M$. For simplicity, we denote $u_{ig} \triangleq u_i(\{g\})$. Without loss of generality, we normalize the utility of each agent such that $\max_{g \in M} u_{ig} = 1$ for each i . An allocation A is a *partition* of the set of goods among the agents; let A_i denote the bundle of goods received by agent i . We want to show that an allocation maximizing the objective $\prod_{i \in N} (1 + u_i(A_i))$ is a $(0, 1)$ -core outcome.

Let A denote an allocation maximizing the smooth Nash welfare objective with $\ell = 1$. We assume without loss of generality that every good is positively valued by at least one agent. Hence, $u_j(A_j) = 0$ must imply $A_j = \emptyset$.

For agents $i, j \in N$ with $A_j \neq \emptyset$ (hence $u_j(A_j) > 0$), and good $g \in A_j$, moving g to A_i should not increase the objective function. Hence, for each $g \in A_j$, we have

$$(1 + u_i(A_i \cup \{g\})) \cdot (1 + u_j(A_j \setminus \{g\})) \leq (1 + u_i(A_i)) \cdot (1 + u_j(A_j)).$$

Using additivity of utilities, this simplifies to

$$\frac{u_{ig}}{1 + u_i(A_i)} \leq \frac{u_{jg}}{1 + u_j(A_j) - u_{jg}} \leq \frac{u_{jg}}{u_j(A_j)}. \quad (6)$$

For every agent $j \in N$ with $A_j \neq \emptyset$ and good $g \in A_j$, define $p_g = u_{jg}/u_j(A_j)$. Abusing the notation a little, for a set $T \subseteq M$ define $p_T = \sum_{g \in T} p_g$. Then, from Equation (6), we have that for all players $i \in N$ and goods $g \in M$,

$$(1 + u_i(A_i)) \cdot p_g \geq u_{ig}. \quad (7)$$

Suppose for contradiction that A is not a $(0, 1)$ -core outcome. Then, there exists a set of agents $S \subseteq N$ and an allocation B of the set of all goods to agents in S such that $(|S|/n) \cdot u_i(B_i) \geq 1 + u_i(A_i)$ for every agent $i \in S$, and at least one inequality is strict. Rearranging the terms and summing over $i \in S$, we have

$$\sum_{i \in S} \frac{u_i(B_i)}{1 + u_i(A_i)} > \sum_{i \in S} \frac{n}{|S|} = n. \quad (8)$$

We now derive a contradiction. For agent $i \in S$, summing Equation (7) over $g \in B_i$, we get

$$(1 + u_i(A_i)) \cdot p_{B_i} \geq u_i(B_i) \Rightarrow \frac{u_i(B_i)}{1 + u_i(A_i)} \leq p_{B_i}.$$

Summing this over $i \in S$, we get

$$\sum_{i \in S} \frac{u_i(B_i)}{1 + u_i(A_i)} \leq \sum_{i \in S} p_{B_i} = \sum_{g \in M} p_g = \sum_{\substack{j \in N \\ A_j \neq \emptyset}} \sum_{\substack{\text{s.t. } g \in A_j}} \frac{u_{jg}}{u_j(A_j)} = \sum_{\substack{j \in N \\ \text{s.t.} \\ A_j \neq \emptyset}} \frac{u_j(A_j)}{u_j(A_j)} \leq n.$$

However, this contradicts Equation (8). □

4 Matching Constraints

We now present the algorithm proving Theorem 2. 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 [19], 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 $F((\mathbf{c} \setminus M(T)) \cup T) - F(\mathbf{c}) \geq \frac{n}{\kappa r}$, 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. The outline of the analysis is similar to the analysis for matroid constraints. First, we show that the algorithm runs in polynomial time. Again, recall that each agent has utility at most m . Thus, $F(\mathbf{c}) = O(n \cdot \ln m)$. Because each improvement increases the objective value by at least $n/(\kappa r)$, the number of iterations is $O(\kappa r \ln m) = O(m^2/\delta)$. Each iteration can be implemented by naïvely going over all $O(m^\kappa)$ subsets of edges of size at most κ , checking if they are valid

augmentations with respect to \mathbf{c} , and whether they improve the objective function by more than $n/(\kappa r)$. The local search therefore runs in polynomial time for constant $\delta > 0$.

Let \mathbf{c} denote the outcome returned by the algorithm. We next show that \mathbf{c} is indeed a $(\delta, 8 + 3\kappa)$ -core outcome. Suppose for contradiction that this is not true. Then, there exist a subset of agents S and a matching \mathbf{c}' such that for all $i \in S$,

$$\frac{|S|}{n} \cdot u_i(\mathbf{c}') \geq (1 + \delta) \cdot u_i(\mathbf{c}) + 8 + 3\kappa \geq (1 + \delta) \cdot (u_i(\mathbf{c}) + 3\kappa + 1),$$

and at least one inequality is strict (the last inequality is because $\delta \in (0, 1]$). Rearranging and summing over all $i \in S$, we obtain

$$\sum_{i \in S} \frac{u_i(\mathbf{c}')}{u_i(\mathbf{c}) + 3\kappa + 1} > (1 + \delta) \cdot \sum_{i \in S} \frac{n}{|S|} = n \cdot (1 + \delta). \quad (9)$$

For $j \in E$, define $w_j = \sum_{i \in N} \frac{u_{ij}}{u_i(\mathbf{c}) + 1}$ and $w'_j = \sum_{i \in N} \frac{u_{ij}}{u_i(\mathbf{c}) + 3\kappa + 1}$. Let $W = \sum_{j \in \mathbf{c}} w_j$, and $W' = \sum_{j \in \mathbf{c}'} w'_j$. It is easy to check that

$$W \leq n \quad \text{and} \quad W' \geq n \cdot (1 + \delta), \quad (10)$$

where the latter follows from Equation (9). Further note that $w_j \geq w'_j$ for all j .

For an augmentation T with respect to \mathbf{c} , define $\text{gain}(T) = \sum_{j \in T} w'_j - \sum_{j \in M(T)} w_j$. The next lemma is a simple generalization of the analysis in [19]; we give the adaptation here for completeness.

Lemma 4. *Assuming weights $w_j \geq w'_j$ for all edges j , for any integer $\kappa \geq 1$ and matchings \mathbf{c} and \mathbf{c}' , there exists a multiset OPT of augmentations with respect to \mathbf{c} such that:*

- For each $T \in OPT$, $T \subseteq \mathbf{c}'$ and $|T| \leq \kappa$;
- $|OPT| \leq \kappa r$; and
- $\sum_{T \in OPT} \text{gain}(T) \geq \kappa \cdot W' - (\kappa + 1) \cdot W$.

Proof. We follow [19] in the construction the multiset OPT of augmentations with respect to \mathbf{c} out of edges in \mathbf{c}' . Let $\mathbf{c} \Delta \mathbf{c}'$ be the symmetric difference of matchings \mathbf{c} and \mathbf{c}' consisting of alternating paths and cycles. For every cycle or path $\mathbf{d} \in \mathbf{c} \Delta \mathbf{c}'$, let $T_{\mathbf{d}}$ be set of edges $\mathbf{d} \cap \mathbf{c}'$. For all $T_{\mathbf{d}}$ with $|T_{\mathbf{d}}| \leq \kappa$, just add $T_{\mathbf{d}}$ to OPT κ times (note that OPT is a multiset, not a set). For $T_{\mathbf{d}}$ with $|T_{\mathbf{d}}| > \kappa$, we break up $T_{\mathbf{d}}$ into multiple smaller augmentations. To do so, index the edges in $T_{\mathbf{d}}$ from 1 to $|T_{\mathbf{d}}|$ and add $|T_{\mathbf{d}}|$ different augmentations to OPT by considering starting at every index in $T_{\mathbf{d}}$ and including the next κ edges in $T_{\mathbf{d}}$ with wrap-around from $|T_{\mathbf{d}}|$ to 1.

Now we must argue that OPT as we have constructed it satisfies the conditions of the lemma. The first point, that $\forall T \in OPT, T \subseteq \mathbf{c}'$ and $|T| \leq \kappa$, follows trivially from the construction. The second point also follows easily from observing that we add κ augmentations to OPT for every $\mathbf{d} \in \mathbf{c} \cap \mathbf{c}'$, and graph G has r vertices.

To see the third point, note that every edge in $\mathbf{c}' \setminus \mathbf{c}$ is contained in at least κ augmentations in OPT . On the other hand, for every edge $e \in \mathbf{c} \setminus \mathbf{c}'$, there are no more than $\kappa + 1$ augmentations $T \in OPT$ such that $e \in M(T)$ (recall $M(T)$ are the edges of \mathbf{c} with a vertex matched under T). This can happen, for example, if T_S happens to be a path of length $\kappa + 1$. Finally, for the edges $j \in \mathbf{c}' \cap \mathbf{c}$, the weight $w'_j \leq w_j$. Putting these facts together, the third point of the lemma follows. \square

Consider the set of augmentations OPT from Lemma 4. For augmentation $T \in OPT$, we have:

$$\begin{aligned}
F((\mathbf{c} \setminus M(T)) \cup T) - F(\mathbf{c}) &= \left(F((\mathbf{c} \setminus M(T)) \cup T) - F(\mathbf{c} \setminus M(T)) \right) - \left(F(\mathbf{c}) - F(\mathbf{c} \setminus M(T)) \right) \\
&\geq \sum_{i \in N} \left(\frac{\sum_{T \in S} u_{ij}}{u_i(\mathbf{c}) + 2\kappa + 1 + \sum_{j \in T} u_{ij}} - \frac{\sum_{j \in M(T)} u_{ij}}{u_i(\mathbf{c}) + 2\kappa + 1 - \sum_{j \in M(T)} u_{ij}} \right) \\
&\geq \sum_{i \in N} \left(\frac{\sum_{j \in T} u_{ij}}{u_i(\mathbf{c}) + 3\kappa + 1} - \frac{\sum_{j \in M(T)} u_{ij}}{u_i(\mathbf{c}) + 1} \right) \\
&= \sum_{j \in T} w'_j - \sum_{j \in M(T)} w_j = \text{gain}(T).
\end{aligned}$$

Here, the second transition holds because $h/(x+h) \leq \ln(x+h) - \ln x \leq h/x$ for all $x \geq 1$ and $h \geq 0$, and the third transition holds due to $|T| \leq \kappa$ and $|M(T)| \leq 2|T| \leq 2\kappa$. Therefore, we have:

$$\begin{aligned}
\sum_{T \in OPT} F((\mathbf{c} \setminus M(T)) \cup T) - F(\mathbf{c}) &\geq \sum_{T \in OPT} \text{gain}(T) \geq \kappa \cdot W' - (\kappa + 1) \cdot W \\
&\geq \kappa \cdot n \cdot (1 + \delta) - (\kappa + 1) \cdot n = n,
\end{aligned}$$

where the second transition follows from Lemma 4, and the third transition follows from Equation (10). Since $|OPT| \leq \kappa r$, there exists an augmentation $T \in OPT$ with $F((\mathbf{c} \setminus M(T)) \cup T) - F(\mathbf{c}) \geq n/\kappa r$, which violates local optimality of \mathbf{c} . This completes the proof of Theorem 2.

Lower Bound. We give a stronger lower bound for matchings than the lower bound for matroids in Lemma 2.

Lemma 5. *A (δ, α) -core outcome is not guaranteed to exist for matching constraints, for any $\delta \geq 0$ and $\alpha < 1$.*

Proof. This example shows that Consider the graph $K_{2,2}$ (the complete bipartite graph with two vertices on each side). This graph has four edges, and two disjoint perfect matchings.

Let there be two agents. Agent 1 has unit utility for the edges of one matching, while agent 2 has unit utility for the edges of the other matching. Any integer outcome gives zero utility to one of these agents. This agent can deviate and obtain utility 2. Hence, for an outcome to be a (δ, α) -core outcome, we need $(1 + \delta) \cdot 0 + \alpha \geq (1/2) \cdot 2$, which is impossible for any $\delta \geq 0$ and $\alpha < 1$. \square

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 3. 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 (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. In words, 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 can be replaced by any positive constant; we choose 1 for simplicity.

We now show an upper bound for R that holds for all instances. Recall from Equation (1) that ρ is the *width* of the instance.

Lemma 6. $R \leq \min(V_{\max}, n, \rho)$, 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:

$$\begin{aligned} & \text{Maximize } \hat{r} && (11) \\ & \sum_{j \in W} u_{ij} w_j \geq V_i \cdot \hat{r} - 1 \quad \forall i \in [n] \\ & A\mathbf{w} \leq \mathbf{b} \\ & w_j \in [0, 1] \quad \forall j \in W \end{aligned}$$

Here, $A\mathbf{w} \leq \mathbf{b}$ 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 3) 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. Fix constant $\delta \in (0, 1)$. Suppose we are given a set of K packing constraints $\mathbf{A}\mathbf{x} \leq \mathbf{b}$ 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; we present Example 3 in Appendix A, which demonstrates this using only one packing constraint. 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.

$$\begin{aligned} & \text{Minimize } \alpha && (12) \\ & \alpha + (1 + \delta) \cdot \sum_{j \in W} u_{ij} w_j \geq U_i^* && \forall i \in [n] \\ & \mathbf{A}\mathbf{w} \leq \mathbf{b} \\ & w_j \in \{0, 1\} && \forall j \in W \\ & \alpha \geq 0 \end{aligned}$$

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 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 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 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 4, where x_j is the fraction of element j chosen.
2. Let \mathbf{y} be an MPF outcome as in Definition 5.
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. 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 4 for details), which does not change our guarantee asymptotically. Further, \mathbf{y} can be computed in polynomial time (Lemma 6). 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.

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, $V_i \geq e^{Q_L} \geq \frac{2RL}{\gamma^3} > 2R$. 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 7. (Proved in Appendix B.2) 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 $\mathbf{E}[X] = (1 - \gamma) \cdot A + \gamma \cdot B$ 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 7 to bound the sizes of F_ℓ for $\ell \in \{0, 1, \dots, L\}$ as follows.

Theorem 6. 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.

Proof. We prove the first and the second part of Theorem 6 separately by considering the heavy groups and the light group in turn. The combined result follows from the union bound.

Case 1: Heavy Groups Consider a heavy group G_ℓ for $0 \leq \ell < L$. Recall that the MPF solution provides utility at least $V_i/(2R)$ to each agent in a heavy group. Hence, we have:

$$\mathbf{E} \left[\widehat{U}_i / (1 - \gamma) \right] = u_i(\mathbf{z}) \geq (1 - \gamma) \cdot U_i^* + \gamma \cdot \frac{V_i}{2R}. \quad (13)$$

The key point is that even if U_i^* is small, the expected utility is at least a term that is proportional to V_i . This will strengthen our application of Chernoff bounds. Using Lemma 7 with $A = U_i^*$ and $B = V_i/(2R)$, we have:

$$\begin{aligned} \Pr \left[\widehat{U}_i < (1 - 3\gamma) \cdot U_i^* \right] &\leq \Pr \left[\frac{\widehat{U}_i}{1 - \gamma} < (1 - 2\gamma) \cdot U_i^* \right] \\ &\leq e^{-\frac{\gamma^3}{4} \frac{V_i}{2R}} \leq e^{-\frac{\gamma^3}{8R} Q_\ell^2} \\ &\leq e^{-Q_\ell \cdot \log L} \leq e^{-(Q_\ell + 2 \log L + \log 6)} \leq \frac{1}{6L^2 e^{Q_\ell}}, \end{aligned} \quad (14)$$

where the second inequality holds because $\log V_i \geq 2 \log Q_\ell$, the third holds because $Q_\ell = \Omega\left(\frac{RL}{\gamma^3}\right)$, and the fourth holds because $L = \omega(1)$.

We are now ready to prove the first part of Theorem 6. Let $\eta_\ell = \frac{1}{6L^2 e^{Q_\ell}}$. Recall that F_ℓ consists of agents in G_ℓ for which $\widehat{U}_i < (1 - 3\gamma) \cdot U_i^*$. Using the linearity of expectation in Equation (14), we have $\mathbf{E}[|F_\ell|] \leq \eta_\ell \cdot |G_\ell|$. By Markov's inequality, $\Pr[|F_\ell| > 3L \cdot \eta_\ell \cdot |G_\ell|] \leq \frac{1}{3L}$. Applying the union bound over the L heavy groups, we have that with probability at least $2/3$,

$$|F_\ell| \leq 3L \cdot \eta_\ell \cdot |G_\ell| = \frac{1}{2Le^{Q_\ell}} \cdot |G_\ell|, \quad \forall \ell \in \{0, 1, \dots, L-1\},$$

which proves the first part of Theorem 6.

Case 2: Light Group For the light group, note that $\log V_i \leq Q_L$. For this group, the MPF solution may not provide any non-trivial guarantee on the utility to the agents. Since the expected utility can now be small, we have to allow additive approximation as well. Recall that F_L consists of agents in G_L for whom $\widehat{U}_i < (1 - 3\gamma) \cdot U_i^*$ as well as $\widehat{U}_i < U_i^* - 4Q_L/\gamma^4$. We again consider two cases.

Case 1. If $U_i^* \leq \frac{4}{\gamma^4} Q_L$, then $\widehat{U}_i \geq U_i^* - \frac{4Q_L}{\gamma^4}$ trivially.

Case 2. Otherwise, $U_i^* \geq \frac{4Q_L}{\gamma^4}$, and using Lemma 7, we have:

$$\Pr\left[\widehat{U}_i < (1 - 3\gamma)U_i^*\right] \leq \Pr\left[\frac{\widehat{U}_i}{1 - \gamma} < (1 - 2\gamma)U_i^*\right] \leq e^{-\frac{\gamma^3}{4}U_i^*} \leq e^{-\frac{\gamma^3}{4} \cdot \frac{4Q_L}{\gamma^4}} \leq \frac{1}{4e^{Q_L}}.$$

It is easy to check that the final transition holds because $\gamma < 1$ is a constant and $Q_L = \omega(1)$.

Note that none of the agents in F_L are in Case 1. Hence, by Markov's inequality, we again have:

$$\Pr\left[|F_L| \geq \frac{1}{2e^{Q_L}} \cdot |G_L|\right] \leq \frac{1}{2},$$

which proves the second part of Theorem 6. □

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 6. 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 7. *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

$$U_i^* \leq \frac{1}{1 - 3\gamma} \cdot \widehat{U}_i + \frac{4Q_L}{\gamma^4}. \quad (15)$$

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 4 only finds an approximate solution whose utilities $\{\widehat{U}_i\}$ with $\sum_{i \in S \cap W} \frac{U_i(\mathbf{h})}{\widehat{U}_i} \leq n(1 + \eta)$ 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 (15) 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:

$$\frac{|S|}{n} \cdot u_i(\mathbf{h}) \leq (1 + 8\gamma) \cdot \widehat{U}_i + \frac{5Q_L}{\gamma^4}. \quad (16)$$

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 (16). 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 6. 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 6, 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 7 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 6 is a (δ, α) -core solution for every $\alpha > \alpha^*$, which completes the proof of Theorem 5.

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 that our algorithms for matroid and matching constraints that globally optimize the smooth Nash welfare objective achieve *exact* rather than approximate Pareto efficiency, in addition to an approximation of the core. An interesting question is whether the same guarantee can be provided (regardless of computation time) for general packing constraints.

Another natural question following our work is to tighten our upper bounds, or to establish matching lower bounds. For instance, we show the existence of a $(0, 2)$ -core outcome for matroid constraints (Theorem 1), 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 Appendix C.

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.

Acknowledgment. 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.

References

- [1] AIRIAU, S., AZIZ, H., CARAGIANNIS, I., KRUGER, J., AND LANG, J. Positional social decision schemes: Fair and efficient randomized voting. Manuscript, 2017.
- [2] AZEVEDO, E. M., AND BUDISH, E. Strategy-proofness in the large. Tech. rep., National Bureau of Economic Research, 2017.
- [3] AZIZ, H., BRILL, M., CONITZER, V., ELKIND, E., FREEMAN, R., AND WALSH, T. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.
- [4] AZIZ, H., ELKIND, E., HUANG, S., LACKNER, M., FERNANDEZ, L. S., AND SKOWRON, P. On the complexity of extended and proportional justified representation. In *AAAI Conference on Artificial Intelligence* (2018).

- [5] AZIZ, H., AND MACKENZIE, S. 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)* (2016), pp. 416–427.
- [6] BENADE, G., NATH, S., PROCACCIA, A. D., AND SHAH, N. Preference elicitation for participatory budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)* (2017), pp. 376–382.
- [7] BUBECK, S. Convex optimization: Algorithms and complexity. *Found. Trends Mach. Learn.* 8, 3-4 (2015), 231–357.
- [8] BUDISH, E. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- [9] CARAGIANNIS, I., KUROKAWA, D., MOULIN, H., PROCACCIA, A. D., SHAH, N., AND WANG, J. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)* (2016), pp. 305–322.
- [10] CONITZER, V., FREEMAN, R., AND SHAH, N. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)* (2017), pp. 629–646.
- [11] FAIN, B., GOEL, A., AND MUNAGALA, K. The core of the participatory budgeting problem. In *Proceedings of the 12th Conference on Web and Internet Economics (WINE)* (2016), pp. 384–399.
- [12] FLUSCHNIK, T., SKOWRON, P., TRIPHAUS, M., AND WILKER, K. Fair knapsack. arXiv:1711.04520, 2017.
- [13] FOLEY, D. K. Lindahl’s solution and the core of an economy with public goods. *Econometrica* 38, 1 (1970), 66–72.
- [14] GANDHI, R., KHULLER, S., PARTHASARATHY, S., AND SRINIVASAN, A. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM* 53, 3 (2006), 324–360.
- [15] GARG, N., KAMBLE, V., GOEL, A., MARN, D., AND MUNAGALA, K. Collaborative optimization for collective decision-making in continuous spaces. In *Proceedings of the 26th International World Wide Web Conference (WWW)* (2017), pp. 617–626.
- [16] GHODSI, A., ZAHARIA, M., HINDMAN, B., KONWINSKI, A., SHENKER, S., AND STOICA, I. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI)* (2011), pp. 24–37.
- [17] GOEL, A., KRISHNASWAMY, A. K., AND SAKSHUWONG, S. Budget aggregation via knapsack voting: Welfare-maximization and strategy-proofness. *Collective Intelligence* (2016).
- [18] GOEL, A., KRISHNASWAMY, A. K., SAKSHUWONG, S., AND AITAMURTO, T. Knapsack voting. *Collective Intelligence* (2015).
- [19] HOUGARDY, S., AND VINKEMEIER, D. E. Approximating weighted matchings in parallel. *Information Processing Letters* 99, 3 (2006), 119 – 123.
- [20] KEARNS, M., NEEL, S., ROTH, A., AND WU, Z. S. Preventing Fairness Gerrymandering: Auditing and Learning for Subgroup Fairness. *ArXiv e-prints* (Nov. 2017).

- [21] KELLY, F. P., MAULLOO, A. K., AND TAN, D. K. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research society* (1998), 237–252.
- [22] KUNG, J. P. S. Basis-exchange properties. In *Theory of Matroids*, W. N., Ed. Cambridge University Press, 1986.
- [23] KUNJIR, M., FAIN, B., K., AND BABU, S. ROBUS: fair cache allocation for data-parallel workloads. In *Proceedings of the 2017 ACM International Conference on Management of Data (SIGMOD)* (2017), pp. 219–234.
- [24] LIPTON, R. J., MARKAKIS, E., MOSSEL, E., AND SABERI, A. On approximately fair allocations of indivisible goods. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)* (2004), pp. 125–131.
- [25] LUZUM, N. Voters to decide fate of durham bond referenda. *The Chronicle* (2016).
- [26] MUENCH, T. J. The core and the lindahl equilibrium of an economy with a public good: An example. *Journal of Economic Theory* 4, 2 (1972), 241 – 255.
- [27] NASH, J. F. The bargaining problem. *Econometrica* 18, 2 (1950), pp. 155–162.
- [28] PARKES, D. C., PROCACCIA, A. D., AND SHAH, N. Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. In *Proceedings of the 13th ACM Conference on Electronic Commerce* (2012), EC '12, pp. 808–825.
- [29] PBP. Participatory budgeting project. Website, 2015.
- [30] PLAUT, B., AND ROUGHGARDEN, T. Almost envy-freeness with general valuations. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2018), pp. 2584–2603.
- [31] PLOTKIN, S. A., SHMOYS, D. B., AND TARDOS, É. Fast approximation algorithms for fractional packing and covering problems. *Math. Oper. Res.* 20, 2 (1995), 257–301.
- [32] PROCACCIA, A. D. Cake cutting: Not just child’s play. *Commun. ACM* 56, 7 (July 2013), 78–87.
- [33] SÁNCHEZ-FERNÁNDEZ, L., ELKIND, E., LACKNER, M., FERNÁNDEZ, N., FISTEUS, J. A., BASANTA VAL, P., AND SKOWRON, P. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)* (2017), pp. 670–676.
- [34] SCARF, H. E. The core of an n person game. *Econometrica* 35, 1 (1967), pp. 50–69.
- [35] SCHUMMER, J. Strategy-proofness versus efficiency on restricted domains of exchange economies. *Social Choice and Welfare* 14 (1997), 47–56.
- [36] THIELE, T. N. Om flerfoldsvalg. In *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger* (1895), pp. 415–441.
- [37] VARIAN, H. R. Two problems in the theory of fairness. *Journal of Public Economics* 5, 3-4 (1976), 249–260.

Appendix

A Impossibility Examples

Example 2. The following example shows the necessity of assuming a large \mathbf{b} in Theorem 3 for general packing constraints. Specifically, the example uses packing constraints $A\mathbf{x} \leq \mathbf{b}$ where $\mathbf{b} = \mathbf{1}$ (and the width is $\rho = 2$), and does not admit a $(\delta, m/4)$ -core outcome for any $\delta > 0$.

Consider a complete bipartite graph $G(L, R, E)$, where $|L| = |R| = m/2$. The vertices are the elements of the ground set W , and the constraints ensure that feasible outcomes are independent sets. There are two agents. Agent 1 has unit utility for each vertex in L , and zero utility for each vertex in R , while agent 2 has unit utility for each vertex in R . A feasible outcome is forced to choose either vertices from L or vertices from R , and hence gives zero utility to at least one agent. But this agent can deviate and choose an outcome with utility $m/2$, which is then scaled down to $m/4$. Hence, no feasible outcome is $(\delta, m/4)$ -core for any $\delta > 0$.

Note that in this example, the welfare-maximizing outcome for a *single* agent is simple to compute, which shows that the non-existence of a good approximation to the core is orthogonal to the computational difficulty of the single-agent welfare maximization problem.

Example 3. Recall that in the KNAPSACK setting, we are given a set of elements of different sizes, and our goal is to select a subset of elements with total size at most a given budget B . We show that for any $\ell > 0$, there exists a KNAPSACK instance in which maximizing the smooth Nash welfare objective $F(\mathbf{c}) = \sum_{i \in N} \ln(\ell + u_i(\mathbf{c}))$ returns an outcome that is not a $(O(m^{1/2-\epsilon}), O(m^{3/4-\epsilon}))$ -core outcome. This is in contrast to Theorem 5, which provides a (δ, α) -core guarantee where δ is constant and α is logarithmic in the number of elements.

Fix $\ell > 0$. Set a large budget $B \geq \ell^4$. There are $m = B^{1/4} + B$ elements, of which $B^{1/4}$ are *large* elements of size $B^{3/4}$ and the remaining B are *small* elements of unit size. There are $n \geq 4B^{1/4} \log(2B)$ agents. Each agent has unit utility for each large element. A subset of αn agents are *special* (we determine α later). These special agents have unit utility for each small element, while the remaining agents have zero utility for the small elements.

The idea is to show that when α is sufficiently small, the smooth Nash welfare objective will choose only the large elements. However, α can still be large enough so that the special agents can deviate, and get a large amount of utility.

Note that maximizing the smooth Nash welfare objective returns a Pareto efficient solution, and hence can be one of two types: it either chooses all large elements (which gives utility $B^{1/4}$ to each agent), or it chooses $B^{1/4} - r$ large elements and $rB^{3/4}$ small elements. For the former to have a larger smooth Nash welfare objective value, we need that for each $1 \leq r \leq B^{1/4}$,

$$\ln(B^{1/4} + \ell) > \alpha \ln(rB^{3/4} + (B^{1/4} - r) + \ell) + (1 - \alpha) \ln(B^{1/4} - r + \ell).$$

This holds true if

$$\ln\left(\frac{B^{1/4} + \ell}{B^{1/4} + \ell - r}\right) > \alpha \ln(rB^{3/4} + (B^{1/4} - r) + \ell).$$

Since $0 < \ell \leq B^{1/4}$, the above is true for each $1 \leq r \leq B^{1/4}$ if

$$\ln\left(\frac{2B^{1/4}}{2B^{1/4} - 1}\right) > \alpha \ln(B^{1/4}B^{3/4} + B^{1/4}).$$

This is true when

$$\ln\left(1 + \frac{1}{2B^{1/4}}\right) \geq \alpha \ln(2B).$$

Since $\ln(1+x) \geq x/2$ for $x \in [0, 1]$, the above holds when:

$$\alpha \leq \frac{1}{4B^{1/4} \ln(2B)}.$$

Let us set $\alpha = \frac{1}{4B^{1/4} \ln(2B)}$. Choosing all large elements maximizes the smooth Nash welfare objective. Since $n \geq 1/\alpha$, there is at least one special agent. The special agents get utility $B^{1/4}$ each. If they deviate and choose all the small elements, they get (scaled down) utility

$$\alpha B = \frac{B}{4B^{1/4} \ln(2B)} = \frac{B^{3/4}}{4 \ln(2B)}.$$

Hence, for the solution to be a (δ, α) -core outcome, we need $(1+\delta) \cdot B^{1/4} + \alpha \geq B^{3/4}/(4 \ln(2B))$. Since $m = \Theta(B)$, this shows that the outcome is not a $(O(m^{1/2-\epsilon}), O(m^{3/4-\epsilon}))$ -core outcome for any constant $\epsilon > 0$, as required.

B Omitted Proofs

B.1 Proof of Theorem 4

The fractional outcome maximizing the Nash welfare objective is the solution of the following program. For simplicity of presentation, we absorb the constraint that $w_j \leq 1$ for each $j \in W$ into the packing constraints.

$$\begin{aligned} \text{Maximize} \quad & \sum_{i \in N} \ln U_i & (17) \\ \sum_{j=1}^m a_{kj} w_j & \leq b_k & \forall k \in [K] \\ U_i & = \sum_{j=1}^m w_j u_{ij} & \forall i \in N \\ w_j & \geq 0 & \forall j \in W \end{aligned}$$

Denote a vector of utilities by $\mathbf{U} = \langle U_1, U_2, \dots, U_n \rangle$, and the polytope of feasible utility vectors by \mathcal{U} . Then, the fractional MNW outcome is obtained by the following maximization.

$$\max_{\mathbf{U} \in \mathcal{U}} \sum_{i \in N} \ln U_i \quad (18)$$

We want to compute a fractional (δ, ϵ) -approximate core outcome in time polynomial in n, V_{\max} , and $\log \frac{1}{\delta \epsilon}$. Assume that \mathcal{U} is a convex polytope of feasible utility vectors. For any $\delta \geq 0, \epsilon > 0$, let $\epsilon' = \epsilon/(1+\delta)$. Define the following objective function. Note that in the absence of the ϵ' term, it would mimic the derivative of the Nash social welfare objective from Program (18).

$$\min_{\mathbf{U} \in \mathcal{U}} Q(\mathbf{U}), \text{ where } Q(\mathbf{U}) = \max_{\mathbf{U}' \in \mathcal{U}} \sum_{i \in N} \frac{U'_i + \epsilon'}{U_i + \epsilon'}. \quad (19)$$

Clearly, $Q(\mathbf{U}) \geq n$ for every \mathbf{U} . Thus, the objective value in Program (19) is at least n . In Section 2.2, we presented an argument showing that the fractional MNW outcome is in the core. A similar argument using the first order optimality condition shows that if $\mathbf{U}^* \in \arg \max_{\mathbf{U} \in \mathcal{U}} \sum_{i \in N} \ln(U_i + \epsilon')$, then

$$\sum_{i \in N} \frac{U_i + \epsilon'}{U_i^* + \epsilon'} \leq n.$$

This implies the optimum of Program (19) is achieved at the fractional outcome maximizing the smooth Nash welfare objective $\sum_{i \in N} \ln(U_i + \epsilon')$, and this optimal value is exactly n .

Next, we turn to efficiently approximating Program (19), and show that if $Q(\mathbf{U}) \leq n(1 + \delta)$, then \mathbf{U} is a (δ, ϵ) -core outcome.

We want to use the Ellipsoid algorithm to approximately minimize the objective function $Q(\mathbf{U})$ over $\mathbf{U} \in \mathcal{U}$ in polynomial time. For this, all we need is that Q is a convex function, its subgradient is efficiently computable, the range of Q and the diameter of \mathcal{U} are exponentially bounded, and polytope \mathcal{U} is efficiently separable [7].

First, we claim that $Q(\mathbf{U})$ is a convex function of \mathbf{U} . To see this, note that for any fixed \mathbf{U}' , U'_i/U_i is convex in U_i . Since the sum and maximum of convex functions is convex, we conclude that $Q(\mathbf{U})$ is also convex.

Second, the subgradient of $Q(\mathbf{U})$ is efficiently computable for every $\mathbf{U} \in \mathcal{U}$. First, we find the $\mathbf{U}' \in \mathcal{U}$ that maximizes $\sum_{i \in N} \frac{U'_i + \epsilon'}{U_i + \epsilon'}$ by solving a linear program. Then, we fix \mathbf{U}' and take the gradient of $\frac{U'_i}{U_i}$ with respect to U_i to obtain a subgradient of $Q(\mathbf{U})$.

Third, note that $U_i \in [0, V_{\max}]$ for each i . Hence, $Q(\mathbf{U}) \leq \frac{n \cdot (V_{\max} + \epsilon')}{\epsilon'}$, which is exponentially bounded in the input size. It is easy to see that the same holds for the diameter of the polytop \mathcal{U} .

Finally, polytope \mathcal{U} is efficiently separable because it is a set of polynomially many linear inequalities.

Hence, we can efficiently obtain a solution $\hat{\mathbf{U}} \in \mathcal{U}$ that satisfies

$$\max_{\mathbf{U}' \in \mathcal{U}} \sum_{i \in N} \frac{U'_i + \epsilon'}{\hat{U}_i + \epsilon'} \leq n + \delta \leq n(1 + \delta).$$

Finally, we show that $\hat{\mathbf{U}}$ must be a (δ, ϵ) -core outcome. Suppose for contradiction that it is not. Then, there exists a subset S of agents and an outcome \mathbf{U}' such that

$$(1 + \delta) \cdot \hat{U}_i + \epsilon \leq \frac{|S|}{n} \cdot U'_i$$

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

$$\sum_{i \in S} \frac{U'_i}{(1 + \delta) \cdot \hat{U}_i + \epsilon} > |S| \cdot \frac{n}{|S|} = n.$$

However, we also have

$$\sum_{i \in S} \frac{U'_i}{(1 + \delta) \cdot \hat{U}_i + \epsilon} \leq \sum_{i \in S} \frac{U'_i + \epsilon'}{(1 + \delta) \cdot (\hat{U}_i + \epsilon')} = \frac{1}{1 + \delta} \sum_{i \in S} \frac{U'_i + \epsilon'}{\hat{U}_i + \epsilon'} \leq n,$$

where the last inequality is due to approximate optimality of $\hat{\mathbf{U}}$. This is a contradiction. Hence, $\hat{\mathbf{U}}$ is a (δ, ϵ) -core outcome.

B.2 Proof of Lemma 7

We first state the standard theorem for Chernoff bounds.

Theorem 8. *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 any $\epsilon \in (0, 1)$, we have:*

$$\Pr[X < (1 - \epsilon)\mathbf{E}[X]] \leq e^{-\frac{\epsilon^2}{2}\mathbf{E}[X]}$$

Equivalently, for any $\eta < \mathbf{E}[X]$,

$$\Pr[X < \mathbf{E}[X] - \eta] \leq e^{-\frac{\eta^2}{2\mathbf{E}[X]}}$$

Lemma 7 follows from considering two cases. First, suppose $(1 - \gamma)A \geq B$.

$$\begin{aligned} \Pr[X < (1 - \gamma)^2 A] &\leq \Pr[X < (1 - \gamma)\mathbf{E}[X]] \leq e^{-\frac{\gamma^2}{2}\mathbf{E}[X]} \\ &\leq e^{-\frac{\gamma^2}{2} \max(\gamma B, (1-\gamma)A)} \leq e^{-\frac{\gamma^3}{2} \max(B, A)} \end{aligned}$$

In the other case, if $(1 - \gamma)A < B$, then $\gamma B \leq \mathbf{E}[X] \leq (1 + \gamma)B$. Then

$$\begin{aligned} \Pr[X < (1 - \gamma)A] &\leq \Pr[X \leq \mathbf{E}[X] - \gamma B] \\ &\leq e^{-\gamma^2 \frac{B^2}{2\mathbf{E}[X]}} \leq e^{-\gamma^2 \frac{B^2}{2(1+\gamma)B}} \\ &\leq e^{-\frac{\gamma^3}{2} B} \leq e^{-\frac{\gamma^3}{2}(1-\gamma)A} \leq e^{-\frac{\gamma^3}{4} A} \end{aligned}$$

C Endowments-Based Core

In this section, we show that our randomized rounding approach to packing problems extends to a slightly different definition of the core, and yields a similar approximation result.

This alternate definition of the core only applies to packing constraints. For simplicity of presentation, we focus on the multi-winner election with general (additive) utilities setting (or knapsack with unit sizes), where n agents have additive utilities over m unit-size elements, and feasible outcomes are subsets of elements of size at most B . So far, we considered a notion of core in which a subset S of agents can deviate and choose a feasible outcome using the entire budget B ; however, their utility is scaled down by $|S|/n$.

A different notion of core is based on scaling the *endowment*. Under this notion, when S deviates, it can choose an outcome with a scaled down budget of $B \cdot |S|/n$, but then its utility is not scaled down. This notion of core has been considered in the context of participatory budgeting [11] and logically implies *proportional representation* of voters in multi-winner elections. This notion builds on the seminal work of [13] on Lindahl equilibrium and its connection to the core.

For $P \leq B$, let $\mathcal{O}(P)$ denote the set of outcomes consisting of at most P elements.

Definition 6. *We say that an outcome \mathbf{c} is a (δ, α) -core outcome if for every subset S of agents and every outcome $\mathbf{c}' \in \mathcal{O}(B \cdot (1 - \delta) \cdot \frac{|S|}{n})$, it is not the case that $u_i(\mathbf{c}') \geq (1 + \delta) \cdot u_i(\mathbf{c}) + \alpha$ for all $i \in S$ and at least one inequality is strict. We refer to a $(0, 0)$ -core outcome simply as a core outcome.*

As shown by [11], it follows directly from the work of [13] that there always exists a fractional core outcome in this setting due to a fixed point argument.² However, it is not known if a fractional core or approximate core outcome can be computed in polynomial time.

More interestingly, it is also an open question whether an integral core outcome always exists for the special case of approval voting, where we assume agents have binary (0 or 1) utilities. It is easy to show that an integral core outcome does not always exist in the more general participatory budgeting setting, in which non-binary utilities and different sized elements are allowed. Consider an example with three elements $\{a, b, c\}$ of size 2 each, a budget of $B = 3$, and three agents with *cyclic* preferences over the elements as follows.

²Note that Theorem 4 does not apply in this setting, since it requires scaling the utility by the size of the coalition.

	a	b	c
u_1	1	0.5	0
u_2	0	1	0.5
u_3	0.5	0	1

An integral outcome \mathbf{c} can only choose a single element. Without loss of generality, suppose $\mathbf{c} = \{a\}$. Then, the set of agents $S = \{2, 3\}$ and outcome $\mathbf{c}' = \{c\}$ show a violation of the core.

We now show the existence of an approximate core solution. We begin with the *fractional core* outcome \mathbf{x} that can be computed using fixed point methods [13, 11]. We use dependent rounding [14] to round \mathbf{x} to \mathbf{X} so that (i) $x_j = \mathbf{E}[X_j]$ for each element j ; (ii) the constraint $\sum_j X_j \leq B$ is preserved; and (iii) $\{X_j\}$ are negatively correlated. Since we do not know if the fractional core outcome can be computed in polynomial time, this algorithm is not necessarily polynomial time, but it yields the following approximation result.

Theorem 9. *For $\delta \in (0, 1]$, there is a (δ, α) -core for multi-winner elections with general (additive) utilities, where $\alpha = O\left(\frac{1}{\delta^4} \log \frac{B}{\delta}\right)$.*

Proof. We only sketch the proof of the upper bound, since it is similar to the argument in Section 5.4.2. Let $\gamma = \frac{\delta}{5}$. Let U_i^* denote the utility agent i receives in the fractional core outcome. We have $U_i^* \in [0, B]$. Let \hat{U}_i be the random variable denoting the utility agent i obtains in the rounded allocation.

Let $L = \frac{2}{\gamma^4} \log \frac{4B}{\gamma}$. First, if $U_i^* \leq L$, then $\hat{U}_i \geq U_i^* - L$ trivially. Otherwise, $U_i^* \geq L$, and using Lemma 7, we have:

$$\Pr \left[\hat{U}_i < (1 - 2\gamma)U_i^* \right] \leq e^{-\frac{\gamma^3}{4}U_i^*} \leq e^{-\frac{\gamma^3}{4}L} \leq \frac{1}{2B}.$$

Let F denote the subset of agents with $\hat{U}_i < \min((1 - 2\gamma)U_i^*, U_i^* - L)$. By Markov's inequality:

$$\Pr \left[|F| \geq \frac{n}{B} \right] \leq \frac{1}{2}.$$

Let $W = N \setminus F$. Suppose set S of agents deviate. We consider two cases:

Case 1. Suppose $|W \cap S| \geq (1 - \gamma)|S|$. Then, consider the agents in $W \cap S$, $P = \frac{|W \cap S|}{n}B \geq (1 - \gamma)\frac{|S|}{n}B \geq (1 - \delta)\frac{|S|}{n}B$ and any allocation $\mathbf{h} \in \mathcal{O}(P)$. By the core condition, there exists $i \in W \cap S$ with $U_i^* \geq U_i(\mathbf{h})$. Since $i \in W$, we have

$$\hat{U}_i \geq \min((1 - 2\gamma)U_i^*, U_i^* - L).$$

This implies $U_i(\mathbf{h}) \leq \hat{U}_i(1 + 5\gamma) + L$.

Case 2. Otherwise, $|S \setminus W| \geq \gamma|S|$. Then

$$|S| \leq \frac{1}{\gamma}|S \cap F| \leq \frac{1}{\gamma}|F| \leq \frac{1}{\gamma} \frac{n}{B}.$$

Thus, if S deviates, their scaled down budget is at most $1/\gamma$. Using this budget, an agent in S can derive utility at most $1/\gamma < \alpha$. Since we give α utility for free to each agent in S under our additive approximation, the approximate core condition is satisfied. \square

The above proof generalizes to arbitrary packing constraints $A\mathbf{x} \leq \mathbf{b}$. In this case, let $\Delta = \max_k \max_j \frac{b_i}{a_{kj}}$. For $P \leq 1$, let $\mathcal{O}(P)$ denote the set of outcomes satisfying $A\mathbf{x} \leq P\mathbf{b}$.

Then, for $\delta \geq 0$ and $\alpha \geq 0$, we say that an outcome \mathbf{c} is a (δ, α) -core outcome if for any $S \subseteq N$ and outcome $\mathbf{c}' \in \mathcal{O}(\frac{t(1-\delta)}{n})$, it is not the case that $u_i(\mathbf{c}') \geq (1 + \delta) \cdot u_i(\mathbf{c}) + \alpha$ and at least one inequality is strict. Generalizing the above proof, it is easy to show the following theorem.

Theorem 10. *For $\delta \in (0, 1]$, there is a (δ, α) -core outcome for general packing problems, where $\alpha = O\left(\frac{1}{\delta^4} \log \frac{\Delta}{\delta}\right)$.*