A constant-factor approximation algorithm for Nash Social Welfare with submodular valuations

Wenzheng Li
Department of Computer Science
Stanford University
Stanford, CA, USA
wzli@stanford.edu

Jan Vondrák
Department of Mathematics
Stanford University
Stanford, CA, USA
jvondrak@stanford.edu

2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS) | 978-1-6654-2055-6/22/$31.00 ©2022 IEEE | DOI: 10.1109/FOCS52979.2021.00012

***Abstract*—We present a 380-approximation algorithm for the Nash Social Welfare problem with submodular valuations. Our algorithm builds on and extends a recent constant-factor approximation for Rado valuations [15].**

***Keywords*-Nash social welfare; approximation algorithms**

I. Introduction

Nash Social Welfare is the following optimization problem.

Nash Social Welfare (NSW): Given m indivisible items and n agents with valuation functions $v_i : 2^{[m]} \to \mathbb{R}_+$, we want to allocate items to the agents, that is find a partition of the m items $(S_1, S_2, \ldots, S_n)$ that maximizes the *geometric average* of the valuations,

$$\mathrm{NSW}(S_1, S_2, \ldots, S_n) = \left(\prod_{i=1}^{n} v_i(S_i)\right)^{1/n}.$$

Among the possible objectives considered in allocation of indivisible goods, it can be viewed as a compromise between Maximum Social Welfare (maximizing the summation $\sum_{i=1}^{n} v_i(S_i)$, which does not take fairness into account), and Max-Min Welfare (maximizing $\min_{1\le i\le n} v_i(S_i)$, which focuses solely on the least satisfied agent and ignores the possible additional benefits to others). The notion of Nash Social Welfare goes back to John Nash's work [20] on bargaining in the 1950s. It also came up independently in the context of competitive equilibria with equal incomes [21] and proportional fairness in networking [17]. An interesting feature of Nash Social Welfare is that the problem is invariant under scaling of the valuations v_i by independent factors λ_i; i.e., each agent can express their preference in a "different currency" and this does not affect the problem.

The difficulty of the problem naturally depends on what class of valuations v_i we consider. Unlike the (additive) Social Welfare Maximization problem, the Nash Social Welfare problem is non-trivial even in the case where the v_i's are *additive*, that is $v_i(S) = \sum_{j\in S} v_{ij}$ where $v_{ij} = v_i(\{j\})$ is agent i's valuation for item j. It is NP-hard in the case of 2 agents with identical additive valuations (by a reduction from the Subset Sum problem), and APX-hard for multiple agents [18]. A constant-factor approximation for the additive case was discovered in a remarkable work by Cole and Gatskelis [11], and subsequently via a very different algorithm by Anari et al. [3]. The algorithm of [11] is based on consideration of market equilibria and market-clearing prices. The algorithm of [3] uses a convex relaxation inspired by Gurvits's work on the permanent of doubly stochastic matrices, which relies on properties of real stable polynomials. Inspired by these exciting breakthroughs, a series of follow-up work has been developed along these two lines [2], [5], [7], [10], [16]. The best approximation factor for additive valuations currently stands at $e^{1/e} \simeq 1.45$ [5].

A particularly compelling question is whether a constant-factor approximation is possible for submodular valuations (where a constant-factor approximation is known for (additive) social welfare maximization [13], [22], and submodular valuations are the largest natural class for which such a result is known, assuming only value-oracle access to the valuations). Some progress has been made for Nash Social Welfare with valuations beyond additive ones: a constant factor for concave piece-wise linear separable utilities [2], and for budget-additive valuations [7], [14]; in fact the approximation factor for budget-additive valuations now matches the $e^{1/e}$ for additive valuations [7]. Recently, [19] designed an algorithm to estimate the optimal value within a factor of $\frac{e^3}{(e-1)^2} \simeq 6.8$ for certain subclasses of submodular valuations, such as coverage and summations of matroid rank functions, by extending the techniques of [3] using stable polynomials. And most recently, [15] designed a constant-factor (772) approximation algorithm for the class of "Rado valuations", which includes matroid rank functions and more generally valuations defined by a certain matching problem with a matroid constraint. [15] presents another significantly different approach to the problem: Instead of market/pricing-inspired techniques or techniques based on stable polynomials, this paper uses a combination of combinatorial matching techniques and a convex programming relaxation.

For general submodular valuations, the best result prior

2575-8454/21/$31.00 ©2021 IEEE
DOI 10.1109/FOCS52979.2021.00012

to this work was an $O(n)$-approximation which also applies to subadditive valuations [4], [16]. However, for subadditive or even fractionally subadditive valuations we cannot expect a constant factor in the value oracle model [4], for the same reasons that this is impossible for the Social Welfare Maximization problem [12]. In the special case of a constant number of agents n with submodular valuations, [16] presents a $(1-1/e-\epsilon)$-approximation for any $\epsilon > 0$; this algorithm uses an extensive enumeration which makes the running time exponential in n.

Our result and techniques:

Theorem 1 (Main Result)**.** *There exists a polynomial-time constant-factor approximation algorithm for the Nash Social Welfare problem with monotone submodular valuation functions, accessible by value queries.*

The approximation factor that we obtain is 380. We made only modest effort to optimize the constant. We believe that the best constant achievable with the techniques of this paper would still be a triple-digit number.

Our techniques can be viewed as a natural extension of the approach in [15]. In hindsight, the strength of the approach of [15] is that it is rather modular and isolates the issue of providing at least some nonzero value to each agent as a separate matching problem. The question then remains how to deal with the remaining items and for this we develop some new techniques. The approach of [15] relies on the existence of a tractable Eisenberg-Gale relaxation with useful polyhedral properties for Rado valuations; this approach might be possibly extended to gross substitutes valuations, but probably not beyond that. The main new components that we introduce are: (i) a new non-convex relaxation of the problem (the Mixed Multilinear Relaxation), (ii) an algorithm to solve it approximately, and (iii) a randomized rounding technique using concentration of submodular functions to obtain an integer solution. We present a more detailed overview at the beginning of Section III.

II. Preliminaries

Nash Social Welfare (NSW): Given a set of m indivisible items $\mathcal{G}$ and a set of n agents $\mathcal{A}$, with valuation functions $v_i : 2^{\mathcal{G}} \to \mathbb{R}_+$ for each $i \in \mathcal{A}$, we want to allocate the items to the agents, that is find a partition $(S_1, S_2, \ldots, S_n)$ of $\mathcal{G}$ in order to maximize the *geometric average* of the valuations,

$$\mathrm{NSW}(S_1, \ldots, S_n) = \left(\prod_{i=1}^{n} v_i(S_i) \right)^{1/n}.$$

Monotone Submodular Functions: Let $\mathcal{G}$ be a finite ground set and $v : 2^{\mathcal{G}} \to \mathbb{R}$.

- v is submodular if for any $S, T \subseteq \mathcal{G}$,

$$v(S) + v(T) \geq v(S \cap T) + v(S \cup T).$$

- v is monotone if $v(S) \leq v(T)$ whenever $S \subseteq T$.

Multilinear Extension: For a set function $v : 2^{\mathcal{G}} \to \mathbb{R}$, we define its multilinear extension $V : [0,1]^{\mathcal{G}} \to \mathbb{R}$ by

$$V(\mathbf{x}) = \sum_{S \subseteq \mathcal{G}} v(S) \prod_{i \in S} x_i \prod_{j \in \mathcal{G} \setminus S} (1 - x_j).$$

The following is well-known and used in prior work (e.g., [6]).

Lemma 2. *Let $V : [0,1]^{\mathcal{G}} \to \mathbb{R}$ be the multilinear extension of a set function $v : 2^{\mathcal{G}} \to \mathbb{R}$. Then*

- *If v is monotone non-decreasing, then V is non-decreasing along any line with direction $\mathbf{d} \geq 0$.*
- *If v is submodular, then V is concave along any line with direction $\mathbf{d} \geq 0$.*

We use the following shorthand notation: For a singleton set $\{j\}$, we write $v(j)$ to denote $v(\{j\})$. For a set S (either containing or not containing j), we write $v(S+j)$ to denote $v(S \cup \{j\})$ and $v(S-j)$ to denote $v(S \setminus \{j\})$. We denote by $\mathbf{1}_S$ the indicator vector of S, i.e. $(\mathbf{1}_S)_j = 1$ if $j \in S$ and 0 otherwise. We also write $\mathbf{1}_j$ instead of $\mathbf{1}_{\{j\}}$ to simplify the notation.

III. Our algorithm and analysis

Algorithm 1 Nash Social Welfare algorithm

procedure NSW($\mathcal{A}, \mathcal{G}, v_1, \ldots, v_n$):
  Find a matching $\tau : \mathcal{A} \to \mathcal{G}$ maximizing $\prod_{i \in \mathcal{A}} v_i(\tau(i))$
  $\mathcal{H} := \tau(\mathcal{A})$, $\mathcal{G}' := \mathcal{G} \setminus \mathcal{H}$, $\mathcal{A}' := \{i \in \mathcal{A} : v_i(\mathcal{G}') > 0\}$
  $\mathbf{y} :=$ **IteratedContinuousGreedy**$(\mathcal{A}', \mathcal{G}', v_1, \ldots, v_n)$
  $(R_1, \ldots, R_n) :=$ **RandomizedRounding**$(\mathbf{y})$
  Find a matching $\sigma : \mathcal{A} \to \mathcal{H}$ maximizing $\prod_{i \in \mathcal{A}} v_i(R_i + \sigma(i))$
  Return $(R_1 + \sigma(1), R_2 + \sigma(2), \ldots, R_n + \sigma(n))$
end procedure

Our algorithm at a high level is described in Algorithm 1. We are strongly inspired by the algorithm of [15] for Rado valuations and follow their high-level structure. We preserve some of the components of their algorithm but replace components which previously relied on special properties of Rado valuations. The new components are: a new relaxation of the Nash Social Welfare problem, and the subroutines **IteratedContinuousGreedy** and **RandomizedRouding**, which are described and analyzed in Sections III-C and III-D, respectively. The analysis can be summarized as follows (with a numbering of phases analogous to [15]).

Phase I: Initial Matching: We find an optimal assignment of 1 item for each agent, i.e. a matching $\tau : \mathcal{A} \to \mathcal{G}$ maximizing $\prod_{i \in \mathcal{A}} v_i(\tau(i))$. This is also the starting point in [15]. $\mathcal{H}$ denotes the items allocated in this matching.

Phase II: Mixed Multilinear Relaxation: We formulate an optimization program which aims to assign the items in $\mathcal{H}$ integrally and the remaining items fractionally under a certain relaxed objective. However, we do not have a concave relaxation at our disposal, such as the Eisenberg-Gale program in [15]; no such tractable relaxation is known for general submodular functions. Instead, we propose a new relaxation involving a product of multilinear polynomials.

$$\begin{aligned} \max \quad & \prod_{i \in \mathcal{A}} V_i(\mathbf{x}_i) && \text{(Mixed-Multilinear)} \\ s.t. \quad & \sum_{i \in \mathcal{A}} x_{ij} \leq 1 && \forall j \in \mathcal{G} \\ & x_{ij} \geq 0 \\ & x_{ij} \in \{0,1\} && \forall i \in \mathcal{A}, j \in \mathcal{H} \end{aligned}$$

Here, $V_i(\mathbf{x}_i) = \sum_{S \subseteq \mathcal{G}} v_i(S) \prod_{j \in S} x_{ij} \prod_{j' \in \mathcal{G} \setminus S} (1 - x_{ij'})$ is the multilinear extension of v_i.

Although the items in $\mathcal{H}$ could be allocated arbitrarily, we will use a matching in the end. Similarly to [15], we prove that this does not hurt the solution significantly. In the next phase, we deal with the question of solving the fractional part of the relaxation.

Phase III: Iterated Continuous Greedy Algorithm: We ignore the items in $\mathcal{H}$ for a moment and try to solve the optimization problem restricted to the item set $\mathcal{G}' = \mathcal{G} \setminus \mathcal{H}$ and the subset of agents $\mathcal{A}'$ who have positive value for these items.

$$\begin{aligned} \max \quad & \prod_{i \in \mathcal{A}'} V_i(\mathbf{y}_i) && \text{(MultilinearProduct)} \\ s.t. \quad & \sum_{i \in \mathcal{A}'} y_{ij} \leq 1 && \forall j \in \mathcal{G}' \\ & y_{ij} \geq 0 \end{aligned}$$

A natural idea is to apply the continuous greedy algorithm of [6]. However, a direct application doesn't work since the objective function is not concave even in nonnegative directions (a product of concave functions is not necessarily concave). We can obtain an objective function concave in nonnegative directions, if we take a logarithm of the objective function: The logarithm of a non-decreasing concave function is non-decreasing concave, and we get a summation instead of a product.

$$\begin{aligned} \max \quad & \sum_{i \in \mathcal{A}'} \log V_i(\mathbf{y}_i) && \text{(LogMultilinear)} \\ s.t. \quad & \sum_{i \in \mathcal{A}'} y_{ij} \leq 1 && \forall j \in \mathcal{G}' \\ & y_{ij} \geq 0 \end{aligned}$$

Nevertheless, the continuous greedy algorithm still doesn't work as such, because it gives a multiplicative approximation; but we require an additive approximation on the logarithmic scale.

Our solution is an iterated version of the continuous greedy algorithm, where we run the continuous greedy algorithm, scale the solution by a factor of $1/2$, and repeat as long as there is some tangible gain. The intuition is that as long as our solution has low value, the continuous greedy process makes progress at a high rate and hence we gain more in the continuous greedy process than what we lose in the scaling step. The output of the iterated continuous greedy algorithm is a solution $\mathbf{y}$ satisfying

$$\sum_{i \in \mathcal{A}'} \frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i)} = O(n)$$

for any feasible solution $\mathbf{y}^*$. This is a stronger guarantee than just approximating the optimum of (LogMultilinear) which will be useful in the analysis.

Phase IV: Randomized Rounding: Our next goal is to round or at least sparsify the fractional solution $\mathbf{y}$. Since our relaxation doesn't have polyhedral properties which were used for sparsification in [15], we resort to a more elementary approach: randomized rounding. We simply allocate each item j to agent i with probability y_{ij}.

Ideally, we would like to argue that the contribution to each agent is strongly concentrated, and thus the value of the assignment is close to the value of the fractional solution. It is known that submodular functions satisfy concentration bounds which can be useful here; the only problem is that the concentration bounds work well only for items with small contributions.

Hence, we partition the items for each agent into "large" and "small": Large items are defined greedily by choosing the maximum marginal profit, as long as the total fractional mass of large items does not exceed some constant $c > 0$. In the analysis, we apply randomized rounding only to the small items. Since their marginal contributions are bounded, we can apply the Efron-Stein inequality and prove that we lose only a constant factor by rounding the small items. The result is a sparsified fractional solution, where only large items are assigned fractionally and their total fractional mass is bounded for each agent.

Phase V: Matching Recombination: The last piece of the puzzle is what to do with large items. Luckily, [15] contains a component which is useful exactly for this purpose. A key lemma in [15] shows that for any fractional solution $\mathbf{y}$ and any matching $\pi : \mathcal{A} \to \mathcal{H}$ (imagine the optimal matching on top of $\mathbf{y}$), there is another matching $\rho : \mathcal{A} \to \mathcal{H}$ such that the value of $(\mathbf{y}, \rho)$ is comparable to the value of $(\mathbf{y}, \pi)$, and for each agent, either the item matched in ρ has a significant value, or there is no item outside of $\mathcal{H}$ which has a significant value. The matching is obtained by an alternating-cycle procedure applied to the matching π and the initial matching τ.

We adapt this lemma and apply it in our setting: After switching to the matching ρ, either the matching item $\rho(i)$ itself provides a constant fraction of agent i's value, or the large items contribute at most a constant fraction of agent i's value. Hence, in both cases we can simply discard the large items in the analysis and lose only a constant factor.

We remark that in the algorithm, we apply randomized rounding to *all* items in $\mathcal{G} \setminus \mathcal{H}$, without distinguishing large and small items. This does not hurt and the algorithm is more natural this way. Also, we do not find the particular matching ρ described here; we simply find the most profitable matching at the end. This provides a solution at least as good as the one we analyze in our proof.

In the following, we describe each phase in detail.

A. Phase I: Initial Matching

First, we solve the Nash Social Welfare problem under the restriction that we only allocate at most one item to each agent. To achieve this, consider the complete bipartite graph between $\mathcal{A}$ and $\mathcal{G}$ and assign an edge weight $\omega_{ij} = \log v_i(j)$ to every edge $(i, j) \in \mathcal{A} \times \mathcal{G}$. We can find an optimal assignment $\tau : \mathcal{A} \to \mathcal{G}$ by computing the maximum-weight matching in this bipartite graph; i.e., $\tau(i)$ is the item matched to agent i. We define $\mathcal{H} = \tau(\mathcal{A})$ to be the set of matched items. We note that each item in the matching has positive value $v_i(\tau(i)) > 0$ for the respective agent, otherwise there is no matching of positive value, which means that $OPT = 0$.

B. Phase II: Mixed Multilinear Relaxation

In this section, we describe our new "Mixed Multilinear" relaxation for the Nash Social Welfare Problem, and a restricted "Matching+Multilinear" version of it, which we show to be within a constant factor of each other. Although these relaxations are new, they are naturally analogous to the relaxations in [15].

Mixed Multilinear Relaxation: For each valuation $v_i : 2^{\mathcal{G}} \to \mathbb{R}_+$, we define its multilinear extension $V_i : [0,1]^{\mathcal{G}} \to \mathbb{R}_+$ as

$$V_i(\mathbf{y}_i) = \sum_{S \subseteq \mathcal{G}} v_i(S) \prod_{j \in S} y_{ij} \prod_{j' \in \mathcal{G}' \setminus S} (1 - y_{ij'}).$$

We propose the following relaxation of the Nash Social Welfare problem.

$$\begin{aligned} \max \quad & \prod_{i \in \mathcal{A}} V_i(\mathbf{x}_i) && \text{(Mixed-Multilinear)} \\ s.t. \quad & \sum_{i \in \mathcal{A}} x_{ij} \leq 1 && \forall j \in \mathcal{G} \\ & x_{ij} \geq 0 && \forall i \in \mathcal{A}, j \in \mathcal{G} \setminus \mathcal{H} \\ & x_{ij} \in \{0,1\} && \forall i \in \mathcal{A}, j \in \mathcal{H} \end{aligned}$$

Note that although $\mathcal{H}$ was chosen by matching one item to each agent, this might not be the case in the optimal solution. Indeed in (Mixed-Multilinear), we allow $\mathcal{H}$ to be allocated arbitrarily; but the assignment cannot be fractional. (If we allowed all items to be assigned fractionally, the relaxation would have an infinite integrality gap, for well-known reasons.) This relaxation is difficult to deal with, because it's hard to find a good assignment of $\mathcal{H}$. Instead, just like in [15], we consider a restricted version of this relaxation, where $\mathcal{H}$ is required to be allocated by a matching.

$$\begin{aligned} \max \quad & \prod_{i \in \mathcal{A}} V_i(\mathbf{y}_i + \mathbf{1}_{\sigma(i)}) && \text{(Matching+Multilinear)} \\ s.t. \quad & \sum_{i \in \mathcal{A}} y_{ij} \leq 1 && \forall j \in \mathcal{G} \setminus \mathcal{H} \\ & y_{ij} \geq 0 && \forall i \in \mathcal{A}, j \in \mathcal{G} \setminus \mathcal{H} \\ & y_{ij} = 0 && \forall i \in \mathcal{A}, j \in \mathcal{H} \\ & \sigma : \mathcal{A} \to \mathcal{H} \text{ is a matching.} \end{aligned}$$

Denote by OPT the optimum value of (Mixed-Multilinear), and by $\text{OPT}_{\mathcal{H}}$ the optimal value of the above program (Matching+Multilinear). Similar to Theorem 3.2 in [15], we have:

Lemma 3.

$$OPT_{\mathcal{H}} \geq \frac{1}{3^{1/3}} OPT.$$

Proof: Consider an optimum solution $\mathbf{x}^*$ of (Mixed-Multilinear), that is $\mathbf{x}^*_i = \mathbf{y}^*_i + \mathbf{1}_{H^*_i}$ where $\mathbf{y}^* \in [0,1]^{\mathcal{A} \times \mathcal{G}}$ is a fractional assignment of the items in $\mathcal{G}'$ and $(H^*_1, \ldots, H^*_n)$ is a partition of $\mathcal{H}$. We construct a feasible solution $(\mathbf{y}^*, \sigma)$ for (Matching+Multilinear), where $\sigma : \mathcal{A} \to \mathcal{H}$ is a matching such that for $H^*_i \neq \emptyset$, $\sigma(i)$ is the most valuable item in H^*_i, and the remaining items in $\mathcal{H}$ are matched arbitrarily to agents such that $H^*_i = \emptyset$.

Let $k_i = |H^*_i|$ be the number of $\mathcal{H}$-items allocated to agent i in the optimal solution. If $k_i > 0$, $\sigma(i)$ is the most valuable of them, and by submodularity $v_i(H^*) \leq k_i v_i(\sigma(i))$. This also implies $V_i(\mathbf{y}^*_i + \mathbf{1}_{H^*_i}) \leq$

$\max\{k_i, 1\} V_i(\mathbf{y}_i^* + \mathbf{1}_{\sigma(i)})$. Hence, we can write

$$\begin{aligned}
\text{OPT} &= \left(\prod_{i\in\mathcal{A}} V_i(\mathbf{y}_i + \mathbf{1}_{H_i^*})\right)^{1/n} \\
&\leq \left(\prod_{i\in\mathcal{A}} \max\{k_i, 1\}\ V_i(\mathbf{y}_i^* + \mathbf{1}_{\sigma(i)})\right)^{1/n} \\
&\leq \left(\prod_{i\in\mathcal{A}} 3^{k_i/3}\ V_i(\mathbf{y}_i^* + \mathbf{1}_{\sigma(i)})\right)^{1/n} \\
&= 3^{1/3}\left(\prod_{i\in\mathcal{A}} V_i(\mathbf{y}_i^* + \mathbf{1}_{\sigma(i)})\right)^{1/n} \\
&\leq 3^{1/3}\ OPT_{\mathcal{H}}
\end{aligned}$$

where we used the AMGM inequality, the fact that $\max\{k, 1\} \leq 3^{k/3}$ for every integer $k \geq 0$, and $\sum_{i=1}^n k_i = n$. ∎

We remark that the factor of $3^{1/3}$ is tight due to the following instance: $|\mathcal{H}| = |\mathcal{G} \setminus \mathcal{H}| = n$, $n/3$ agents have the valuation $v(S) = |S \cap \mathcal{H}|$, and the remaining $2n/3$ agents have the valuation $v'(S) = \min\{|S|, 1\}$. The optimal Nash Social Welfare is $3^{1/3}$, since $n/3$ agents can get value 3 from 3 items of $\mathcal{H}$ each, and the remaining agents get value 1 from items in $\mathcal{G} \setminus \mathcal{H}$. If $\mathcal{H}$ is allocated as a matching, we get Nash Social Welfare 1, since each agent receives value 1.

C. Phase III: The Iterated Continuous Greedy Algorithm

In this section, we describe the details of Phase III where we aim to find a fractional solution of our (LogMultilinear) relaxation of Nash Social Welfare. We do this for a subset of items $\mathcal{G}' = \mathcal{G} \setminus \mathcal{H}$, and a subset of agents $\mathcal{A}'$ who have positive value for these items.

$$\begin{aligned}
\max \quad & \frac{1}{n}\sum_{i\in\mathcal{A}'} \log V_i(\mathbf{y}_i) && \text{(LogMultilinear)} \\
s.t. \quad & \sum_{i\in\mathcal{A}'} y_{ij} \leq 1 && \forall j \in \mathcal{G}' \\
& y_{ij} = 0 && \forall i \in \mathcal{A}', j \in \mathcal{H} \\
& \mathbf{y} \geq 0
\end{aligned}$$

We recall that

$$V_i(\mathbf{y}_i) = \sum_{S\subseteq\mathcal{G}} v_i(S) \prod_{j\in S} y_{ij} \prod_{j'\in\mathcal{G}\setminus S} (1 - y_{ij'})$$

is the multilinear extension of v_i. In this section we assume that the vector $\mathbf{y}_i$ always has 0 in coordinates indexed by $j \in \mathcal{H}$, so effectively we are working with vectors in $[0, 1]^{\mathcal{G}'}$.

We design a variant of the continuous greedy algorithm which approximates the optimal solution within an additive error of 1.

The Iterated Continuous Greedy Algorithm:

1) Start with a feasible solution $\mathbf{y}^{(0)}$, $y_{ij}^{(0)} = \frac{1}{n}$ for each $i \in \mathcal{A}'$ and $j \in \mathcal{G}'$.
2) For $r = 1, 2, \ldots$, given a feasible solution $\mathbf{y}^{(r-1)}$, initiate $\mathbf{y}(\frac{1}{2}) = \frac{1}{2}\mathbf{y}^{(r-1)}$ and run the following continuous greedy algorithm:
   - Let $\mathbf{z}(t)$ be a feasible solution (satisfying $\mathbf{z} \geq 0$ and $\sum_i z_{ij} \leq 1$ for each j) which maximizes the linear objective function
   $$\sum_{i\in\mathcal{A}'} \frac{\mathbf{z}_i \cdot \nabla V_i(\mathbf{y}_i(t))}{V_i(\mathbf{y}_i(t))}.$$
   - Evolve the solution $\mathbf{y}(t)$ according to the equation
   $$\frac{d}{dt}\mathbf{y}(t) = \mathbf{z}(t),$$
   for $t \in [\frac{1}{2}, 1]$.
3) Set $\mathbf{y}^{(r)} = \mathbf{y}(1)$, the solution obtained in this iteration.
4) If
   $$\frac{1}{n}\sum_{i\in\mathcal{A}'} \log V_i(\mathbf{y}^{(r)}) \geq \frac{1}{n}\sum_{i\in\mathcal{A}'} \log V_i(\mathbf{y}^{(r-1)}) + \frac{1}{8},$$
   let $r \leftarrow r + 1$ and repeat.
5) Otherwise, return $\mathbf{y}^{(r)}$.

Theorem 4. *Let $\mathbf{y}^*$ denote any feasible solution of the optimization program (LogMultilinear). Assuming that $v_i(\mathcal{G}') > 0$ and v_i is monotone submodular for each $i \in \mathcal{A}'$, the Iterated Continuous Greedy algorithm terminates in $O(\log n)$ iterations and returns a feasible solution $\mathbf{y}$ for (LogMultilinear) such that*

$$\frac{1}{n}\sum_{i\in\mathcal{A}'} \frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i)} \leq e.$$

We note that by concavity of the logarithm, the conclusion also implies $\frac{1}{n}\sum_{i\in\mathcal{A}'} \log \frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i)} \leq 1$, i.e. our solution approximates the optimum of (LogMultilinear) within an additive error of 1. The statement in the theorem is stronger and more convenient, though, which we will use later in several places.

Proof: As a starting point, we have $y_{ij}^{(0)} = \frac{1}{n}$. By concavity of V_i in positive directions, we have the simple bound $V_i(\mathbf{y}_i^{(0)}) \geq \frac{1}{n} V_i(\mathbf{1})$. Hence, $\frac{1}{n}\sum_{i\in\mathcal{A}'} \log V_i(\mathbf{y}_i^{(0)}) \geq \frac{1}{n}\sum_{i\in\mathcal{A}'} \log V_i(\mathbf{1}) - \log n$. Now we apply the continuous greedy algorithm as above, and we iterate as long as after each iteration we have $\frac{1}{n}\sum_{i\in\mathcal{A}'} \log V_i(\mathbf{y}^{(r)}) \geq \frac{1}{n}\sum_{i\in\mathcal{A}'} \log V_i(\mathbf{y}^{(r-1)}) + \frac{1}{8}$. Since for any feasible solution, $\sum_{i\in\mathcal{A}'} \log V_i(\mathbf{y}^{(r)}) \leq \sum_{i\in\mathcal{A}'} \log V_i(\mathbf{1})$, this means that we cannot iterate more than $O(\log n)$ times. It remains to prove that the solution satisfies the claimed inequality.

To prove this, assume at any time t that

$$\sum_{i\in\mathcal{A}'} \frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i)} > en. \qquad (*)$$

A possible direction for the continuous greedy algorithm to pursue is always $\mathbf{z} = \mathbf{y}^*$. For this direction, we obtain

$$\begin{aligned}
&\sum_{i \in \mathcal{A}'} \frac{\mathbf{y}_i^* \cdot \nabla V_i(\mathbf{y}_i(t))}{V_i(\mathbf{y}_i(t))} \\
&\geq \sum_{i \in \mathcal{A}'} \frac{V_i(\mathbf{y}_i^*) - V_i(\mathbf{y}_i(t))}{V_i(\mathbf{y}_i(t))} \\
&= \sum_{i \in \mathcal{A}'} \left(\frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i(t))} - 1 \right) > (e-1)n
\end{aligned}$$

using the monotonicity and concavity of V_i in nonnegative directions in the first inequality, and our assumption (*) in the second inequality. Since the continuous greedy algorithm chooses a direction $\mathbf{z}_i(t)$ by optimizing the expression $\sum_{i \in \mathcal{A}'} \frac{\mathbf{z}_i \cdot \nabla V_i(\mathbf{y}_i(t))}{V_i(\mathbf{y}_i(t))}$, we obtain the same bound for the greedy direction $\mathbf{z}_i(t)$, and finally by the chain rule we have

$$\begin{aligned}
&\frac{d}{dt} \sum_{i \in \mathcal{A}'} \log V_i(\mathbf{y}_i(t)) \\
&= \sum_{i \in \mathcal{A}'} \frac{1}{V_i(\mathbf{y}_i(t))} \nabla V_i(\mathbf{y}_i(t)) \cdot \frac{d\mathbf{y}_i}{dt} \\
&= \sum_{i \in \mathcal{A}'} \frac{\mathbf{z}_i(t) \cdot \nabla V_i(\mathbf{y}_i(t))}{V_i(\mathbf{y}_i(t))} > (e-1)n.
\end{aligned}$$

Hence the rate of increase in $\sum_{i \in \mathcal{A}'} \log V_i(\mathbf{y}_i(t))$ is at least $(e-1)n$ as long as (*) is satisfied.

Each iteration starts by scaling the previous solution by a factor of $\frac{1}{2}$ and then running continuous greedy for t between $\frac{1}{2}$ and 1. Again by concavity, we have $V_i(\frac{1}{2}\mathbf{y}^{(r-1)}) \geq \frac{1}{2} V_i(\mathbf{y}^{(r-1)})$. By integration over the course of the continuous greedy process, we obtain

$$\begin{aligned}
&\sum_{i \in \mathcal{A}'} \log V_i(\mathbf{y}_i^{(r)}) \\
&= \sum_{i \in \mathcal{A}'} \log V_i \left(\frac{1}{2} \mathbf{y}_i^{(r-1)} \right) + \int_{1/2}^{1} \frac{d}{dt} \sum_{i \in \mathcal{A}'} \log V_i(\mathbf{y}_i(t)) dt \\
&\geq \sum_{i \in \mathcal{A}'} \log \left(\frac{1}{2} V_i(\mathbf{y}_i^{(r-1)}) \right) + \int_{1/2}^{1} (e-1)n \; dt \\
&= \sum_{i \in \mathcal{A}'} \log V_i(\mathbf{y}_i^{(r-1)}) + \left(\frac{e-1}{2} - \log 2 \right) n.
\end{aligned}$$

We note that all logarithms here are natural and $\frac{e-1}{2} - \log 2 > \frac{1}{8}$. Hence we gain at least $\frac{1}{8}n$ in each iteration as long as (*) is satisfied, and we terminate otherwise. ■

Discretization: As in the original continuous greedy algorithm [6], we need to discretize the continuous process to obtain an actual polynomial-time algorithm. This can be done using standard methods.

First, for any given $\mathbf{y}_i(t)$, we can estimate by random sampling

$$\left. \frac{\partial V_i}{\partial y_j} \right|_{\mathbf{y}_i(t)} = \mathbb{E}[v_i(R_i(t) + j) - v_i(R_i(t) - j)]$$

where $R_i(t)$ is a random set containing each item j independently with probability $y_{ij}(t)$. Since $v_i(R_i(t) + j) - v_i(R_i(t) - j) \in [0, v_i(\{j\})]$, using $poly(m, n)$ samples we can obtain estimates ω_{ij} of $\frac{\partial V_i}{\partial y_j}$ within an error of $\frac{v_i(\{j\})}{poly(m,n)}$ with high probability.

Then we find a direction $\mathbf{z}(t)$ by solving the linear programming problem

$$\max \left\{ \sum_{i \in \mathcal{A}'} \frac{1}{V_i(\mathbf{y}_i(t))} \sum_{j \in \mathcal{G}'} \omega_{ij} z_{ij} : z_{ij} \geq 0, \sum_i z_{ij} \leq 1 \; \forall j \right\}$$

(using ω_{ij} in place of $\frac{\partial V_i}{\partial y_j}$). If the estimates ω_{ij} are correct up to an error of $\frac{v_i(\{j\})}{poly(m,n)}$, the optimum is correct up to a relative error of $\frac{1}{poly(m,n)}$. Note that $V_i(\mathbf{y}_i(t)) \geq \frac{1}{poly(m,n)} \sum_{i \in \mathcal{G}'} v_i(\{j\})$ since this is true for the initial solution $\mathbf{y}^{(0)}$ and the value can only decrease $O(\log n)$ times by a factor of 2; apart from that it increases.

Then we make a step of size $\delta = \frac{1}{poly(m,n)}$, where we set $\mathbf{y}(t + \delta) = \mathbf{y}(t) + \delta \cdot \mathbf{z}(t)$. The guarantee we claim here is that

$$\begin{aligned}
&\sum_{i \in \mathcal{A}'} \log V_i(\mathbf{y}_i(t + \delta)) \\
&\geq \sum_{i \in \mathcal{A}'} \log V_i(\mathbf{y}_i(t)) \\
&+ \delta \left(1 - \frac{1}{poly(m,n)} \right) \sum_{i \in \mathcal{A}'} \frac{\mathbf{y}_i^* \cdot \nabla V_i(\mathbf{y}_i(t))}{V_i(\mathbf{y}_i(t))}.
\end{aligned}$$

This is true because we find the optimum of the linear programming problem within a $\frac{1}{poly(m,n)}$ relative error, and also the values $V_i(\mathbf{y}_i)$ and the partial derivatives $\frac{\partial V_i}{\partial y_j}$ can change only by a factor of $1 \pm \frac{1}{poly(m,n)}$ between $\mathbf{y}_i(t)$ and $\mathbf{y}_i(t+\delta)$, as long as $t \leq 0.99$ (since $V_i(\mathbf{y}_i)$ and $\frac{\partial V_i}{\partial y_j}$ are nonnegative and linear in each coordinate separately). Hence, we can mimic the continuous analysis for $t \in [0.5, 0.99]$ within a $\frac{1}{poly(m,n)}$ relative error at every step, and we lose a factor of $49/50$ by ignoring the improvement between $[0.99, 1]$. These errors are easily absorbed for example in the gap between $\frac{e-1}{2} - \log 2$ and $\frac{1}{8}$ which we ignore above. So the theorem still holds for the discretized algorithm, with high probability.

D. Phase IV: Randomized Rounding

In this section, our goal is to round the fractional solution $\mathbf{y}$ from Section III-C. In the actual algorithm, we use the following simple randomized rounding procedure.

RandomizedRounding($\mathbf{y}$):

1) For each item $j \in \mathcal{G}'$ independently, select $Z_j \in \{0, 1, \ldots, n\}$ where $Z_j = i$ with probability y_{ij}, or $Z_j = 0$ with probability $1 - \sum_{i \in \mathcal{A}'} y_{ij}$.
2) Define $R_i = \{j \in \mathcal{G} \setminus \mathcal{H} : Z_j = i\}$.
3) Return $(R_1, \ldots, R_n)$.

However, in the analysis we will proceed more carefully, separating the contributions of "large" and "small" items. We first define what we mean by "large" and "small". For any $\mathbf{y} \in [0,1]^{\mathcal{G}}$ and $S \subseteq \mathcal{G}$, define vector $\mathbf{y}^{(S)}$ to be the vector obtained by setting all the coordinates not in S to 0.

$$y_i^{(S)} = \begin{cases} y_i & i \in S \\ 0 & i \notin S. \end{cases}$$

For each agent $i \in \mathcal{A}'$, we define the set L_i of "large items" as follows, for a given constant $c > 0$. Let us assume in the following that $\sum_{j \in \mathcal{G}'} y_{ij} \geq c$ for every agent i. This is without loss of generality, since we can always extend the instance with dummy items of value 0, which can be allocated fractionally to any agent and it doesn't change the outcome of our algorithm in any way.

FindLargeSet(i, $\mathbf{y}$*):*

1) Start with an empty set at time 0, $L_i^{(0)} = \emptyset$.
2) At time $t \geq 1$, add the item with the largest marginal value to $L_i^{(t-1)}$. More specifically, let $L_i^{(t)} = L_i^{(t-1)} \cup \{j_t\}$ where
$$j_t = \operatorname*{argmax}_{j \in \mathcal{G}' \setminus L_i^{(t-1)}} \left(V_i(\mathbf{y}_i^{(L_i^{(t-1)})} + \mathbf{1}_j) - V_i(\mathbf{y}_i^{(L_i^{(t-1)})}) \right)$$
3) As long as $\sum_{t'=1}^{t} y_{ij_{t'}} < c$ and $\mathcal{G}' \setminus L_i^{(t)} \neq \emptyset$, let $t \leftarrow t+1$ and repeat step 2.
4) Return $L_i := L_i^{(t)}$.

We have two simple corollaries for the set L_i.

- For any agent $i \in \mathcal{A}'$,
$$c \leq \sum_{j \in L_i} y_{ij} < c+1,$$
- For any $i \in \mathcal{A}'$, $j \in \mathcal{G}' \setminus L_i$,
$$V_i(\mathbf{y}_i^{(L_i)} + \mathbf{1}_j) - V_i(\mathbf{y}_i^{(L_i)}) \leq \frac{1}{c} V_i(\mathbf{y}_i^{(L_i)}).$$

The first property follows from the stopping rule (including our assumption that each agent gets $\sum_{j \in \mathcal{G}'} y_{ij} \geq c$ in the fractional solution). As for the second one, if the marginal value is 0 for any $j \in \mathcal{G}' \setminus L_i^{(t-1)}$, it is trivially true. Otherwise, consider any item j in L_i that we did not include in the procedure; (by submodularity) in every step we included an item j_t of marginal value

$$V_i(\mathbf{y}_i^{(L_i^{(t-1)})} + \mathbf{1}_{j_t}) - V_i(\mathbf{y}_i^{(L_i^{(t-1)})}) \geq V_i(\mathbf{y}_i^{(L_i)} + \mathbf{1}_j) - V_i(\mathbf{y}_i^{(L_i)})$$

and by multilinearity the total contribution of the included items is

$$\begin{aligned} V_i(\mathbf{y}_i^{(L_i)}) &= \sum_{t=1}^{|L_i|} y_{ij_t} (V_i(\mathbf{y}_i^{(L_i^{(t-1)})} + \mathbf{1}_{j_t}) - V_i(\mathbf{y}_i^{(L_i^{(t-1)})})) \\ &\geq c(V_i(\mathbf{y}_i^{(L_i)} + \mathbf{1}_j) - V_i(\mathbf{y}_i^{(L_i)})). \end{aligned}$$

Now we can describe our modified rounding procedure. We note that this procedure is used only in the analysis.

RestrictedRandomizedRounding($\mathbf{y}$*):*

1) Compute the set L_i (specified above) for each agent $i \in \mathcal{A}'$.
2) For each item $j \in \mathcal{G}'$, assign j to a random player according to y_{ij}:
Let $Z_j = i$ with probability y_{ij}, or $Z_j = 0$ with probability $1 - \sum_{i \in \mathcal{A}'} y_{ij}$.
For each $i \in \mathcal{A}'$, let $S_i = \{j \in \mathcal{G}' \setminus L_i : Z_j = i\}$ and $\mathbf{y}_i^{(s)} = \mathbf{y}_i^{(L_i)} + \mathbf{1}_{S_i}$.
3) Return $\mathbf{y}^{(s)}$.

Note that only "small items" are included in the sets $S_1, \ldots, S_n$, and large items are still assigned fractionally in $\mathbf{y}^{(s)}$. Thus the solution $\mathbf{y}^{(s)}$ can be viewed as "sparsified" rather than rounded. We note that the notion of sparsity here is in terms of the summation of fractional variables ($\sum_{i \in \mathcal{A}'} \sum_{j \in L_i} y_{ij}^{(s)} < (c+1)n$) rather than the size of the support of $\mathbf{y}^{(s)}$.

The notion of large/small is agent-specific, so $\mathbf{y}^{(s)}$ might not even be a feasible solution; an item could be allocated fully as a small item and still fractionally as a large item for other agents. We will show at the end that large items can be in fact discarded. However, for now we analyze the value of $\mathbf{y}^{(s)}$.

Lemma 5. *Suppose that*

$$\frac{1}{n} \sum_{i \in \mathcal{A}'} \frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i)} \leq \alpha.$$

Then with probability $\Omega(\epsilon)$*, the solution* $\mathbf{y}^{(s)}$ *obtained by* **RestrictedRandomizedRounding(**$\mathbf{y}$**)** *with parameter* $c > 0$ *satisfies*

$$\frac{1}{n} \sum_{i \in \mathcal{A}'} \frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i^{(s)})} \leq (1+\epsilon)(2+4/c)\alpha.$$

Proof: Using the notation from **RestrictedRandomizedRounding(**$\mathbf{y}$**)**, for every $i \in \mathcal{A}'$, we define a monotone submodular function $u_i : 2^{\mathcal{G}' \setminus L_i} \to \mathbb{R}$, where $u_i(S) = V_i(\mathbf{y}_i^{(L_i)} + \mathbf{1}_S)$. Recall that $\mathbf{y}_i^{(s)} = \mathbf{y}_i^{(L_i)} + \mathbf{1}_{S_i}$; that is, $V_i(\mathbf{y}_i^{(s)}) = u_i(S_i)$. The sets $S_1, \ldots, S_n$ are determined by the random variables $(Z_j : j \in \mathcal{G}')$. Our goal is to upper-bound

$$\begin{aligned} V(\mathcal{Z}) = V(Z_j : j \in \mathcal{G}') &= \frac{1}{n} \sum_{i \in \mathcal{A}'} \frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i^{(s)})} \\ &= \frac{1}{n} \sum_{i \in \mathcal{A}'} \frac{V_i(\mathbf{y}_i^*)}{u_i(S_i)}. \end{aligned}$$

By the definition of L_i and by submodularity, we know that for any $i \in \mathcal{A}'$, $j \in \mathcal{G}' \setminus L_i$ and $S \subseteq \mathcal{G}' \setminus L_i$,

$$0 \leq u_i(S \cup \{j\}) - u_i(S) \leq \frac{V_i(\mathbf{y}_i^{(L_i)})}{c} = \frac{u_i(\emptyset)}{c}.$$

Since $u_i(S_i)$ is a function of the independent random variables $(Z_j : j \in \mathcal{G}')$, by the Efron-Stein inequality, we have

$$\begin{aligned}\mathrm{Var}[u_i(S_i)] &\le \mathbb{E}\left[\sum_{j\in\mathcal{G}'}\left(u_i(S_i) - \min_{Z_j} u_i(S_i)\right)^2\right] \\ &= \mathbb{E}\left[\sum_{j\in S_i}\left(u_i(S_i) - u_i(S_i\backslash\{j\})\right)^2\right] \\ &\le \frac{u_i(\emptyset)}{c}\cdot\mathbb{E}\left[\sum_{j\in S_i}\left(u_i(S_i) - u_i(S_i\backslash\{j\})\right)\right] \\ &\le \frac{u_i(\emptyset)}{c}\cdot\mathbb{E}[u_i(S_i)] = \frac{u_i(\emptyset)}{c}\cdot V_i(\mathbf{y}_i)\end{aligned}$$

where we used the submodularity of u_i in the last inequality. By Chebyshev's inequality, we have

$$\Pr\left[u_i(S_i) \le \frac{V_i(\mathbf{y}_i)}{2}\right] \le \frac{\mathrm{Var}[u_i(S_i)]}{(V_i(\mathbf{y}_i)/2)^2} \le \frac{4u_i(\emptyset)}{c\,V_i(\mathbf{y}_i)}.$$

Therefore,

$$\begin{aligned}\mathbb{E}\left[\frac{V_i(\mathbf{y}_i)}{u_i(S_i)}\right] &\le \frac{V_i(\mathbf{y}_i)}{V_i(\mathbf{y}_i)/2} + \frac{V_i(\mathbf{y}_i)}{u_i(\emptyset)}\cdot\Pr\left[u_i(S_i) \le \frac{V_i(\mathbf{y}_i)}{2}\right] \\ &\le 2 + \frac{4}{c}.\end{aligned}$$

Combining this with $\frac{1}{n}\sum_{i\in\mathcal{A}'}\frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i)} \le \alpha$, we can write

$$\mathbb{E}[V(\mathcal{Z})] = \frac{1}{n}\sum_{i\in\mathcal{A}'}\frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i)}\mathbb{E}\left[\frac{V_i(\mathbf{y}_i)}{u_i(S_i)}\right] \le (2+4/c)\alpha.$$

By Markov's inequality, we conclude that with probability $\Omega(\epsilon)$, $V(\mathcal{Z}) \le (1+\epsilon)(2+4/c)\alpha$. ∎

E. Phase V: Matching recombination

Now we have a fractional solution $\mathbf{y}^{(s)}$ with good properties; however, we ignored the fact that $\mathcal{H}$ should be also allocated. Our goal in this section is to prove that there exists a matching which works well with our fractional solution $\mathbf{y}^{(s)}$, and at the same time it has additional properties which allow us round the large items (or in fact discard them!) and still obtain a good value of Nash Social Welfare.

We proceed very much as in [15]. First, we prove that there exists a matching σ which obtains a good value together with $\mathbf{y}^{(s)}$.

Matching extension: Here we show that there exists a matching $\sigma : \mathcal{A} \to \mathcal{H}$ which complements well the fractional solution $\mathbf{y}^{(s)}$.

Lemma 6. *Let $\mathbf{x}^*$ be the optimal solution of (MixedMultilinear), i.e. $\mathbf{x}_i^* = \mathbf{y}_i^* + \mathbf{1}_{H_i^*}$ where $\mathbf{y}^* \in [0,1]^{\mathcal{A}\times\mathcal{G}}$ is a feasible solution of (LogMultilinear) and $(H_1^*, \dots, H_n^*)$ is a partition of $\mathcal{H}$. Let $\mathbf{y}' \in [0,1]^{\mathcal{A}\times\mathcal{G}}$ be an arbitrary fractional solution, satisfying*

$$\frac{1}{n}\sum_{i\in\mathcal{A}'}\frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i')} \le \beta$$

and $\mathbf{y}_i' = 0$ for $i \notin \mathcal{A}'$. Then there is a matching $\pi : \mathcal{A} \to \mathcal{H}$ such that

$$\begin{aligned}NSW(\mathbf{y}', \pi) &= \left(\prod_{i\in\mathcal{A}} V_i(\mathbf{y}_i' + \mathbf{1}_{\pi(i)})\right)^{1/n} \\ &\ge \frac{1}{\beta+1}\left(\prod_{i\in\mathcal{A}} V_i(\mathbf{x}_i^*)\right)^{1/n} \\ &= \frac{1}{\beta+1}\,OPT.\end{aligned}$$

Proof: Suppose that agent i receives $k_i = |H_i^*|$ items from $\mathcal{H}$ in the optimal solution, and let $\pi(i) \in H_i^*$ be the most valuable item in H_i^* (as a singleton). We extend this to a matching $\pi : \mathcal{A} \to \mathcal{H}$, by allocating any remaining items arbitrarily to agents such that $H_i^* = \emptyset$. By the AMGM inequality, we can write

$$\begin{aligned}\frac{OPT}{NSW(\mathbf{y},\pi)} &= \left(\prod_{i\in\mathcal{A}}\frac{V_i(\mathbf{x}_i^*)}{V_i(\mathbf{y}_i + \mathbf{1}_{\pi(i)})}\right)^{1/n} \\ &\le \frac{1}{n}\sum_{i\in\mathcal{A}}\frac{V_i(\mathbf{x}_i^*)}{V_i(\mathbf{y}_i + \mathbf{1}_{\pi(i)})}\end{aligned}$$

By submodularity, we have $V_i(\mathbf{x}_i^*) = V_i(\mathbf{y}_i^* + \mathbf{1}_{H_i^*}) \le V_i(\mathbf{y}_i^*) + k_i V_i(\mathbf{1}_{\pi(i)})$. Thus, we obtain

$$\begin{aligned}\frac{OPT}{NSW(\mathbf{y},\pi)} &\le \frac{1}{n}\sum_{i\in\mathcal{A}}\frac{V_i(\mathbf{y}_i^*) + k_i V_i(\mathbf{1}_{\pi(i)})}{V_i(\mathbf{y}_i + \mathbf{1}_{\pi(i)})} \\ &\le \frac{1}{n}\left(\sum_{i\in\mathcal{A}'}\frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i + \mathbf{1}_{\pi(i)})} + \sum_{i\in\mathcal{A}} k_i\right) \\ &\le \frac{1}{n}\left(\sum_{i\in\mathcal{A}'}\frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i)} + \sum_{i\in\mathcal{A}} k_i\right)\end{aligned}$$

using monotonicity of V_i in the denominator. Note that $V_i(\mathbf{y}_i^*) = 0$ and $v_i(\pi(i)) > 0$ for every agent $i \notin \mathcal{A}'$, because these agents do not derive any value from $\mathcal{G}' = \mathcal{G}\backslash\mathcal{H}$ and hence $v_i(H_i^*) > 0$ for these agents; that's why we can switch to $\mathcal{A}'$ in the first summation. Finally, using the assumption $\frac{1}{n}\sum_{i\in\mathcal{A}'}\frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i)} \le \beta$ and the fact that $\frac{1}{n}\sum_{i\in\mathcal{A}} k_i = \frac{1}{n}\sum_{i\in\mathcal{A}}|H_i^*| = 1$, we obtain

$$\frac{OPT}{NSW(\mathbf{y},\pi)} \le \beta + 1.$$

∎

Corollary 7. *The fractional solution $\mathbf{y}^{(s)}$ =* **RestrictedRandomizedRounding**$(\mathbf{y})$ *satisfies with constant*

probability

$$\max_{\pi} NSW(\mathbf{y}^{(s)}, \pi) \geq \frac{1}{7+12/c} OPT.$$

Proof: Since $\mathbf{y}$ from the Iterated Continuous Greedy algorithm satisfies $\frac{1}{n}\sum_{i\in\mathcal{A}'} \frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i)} \leq e$, we apply Lemma 5 with $\alpha = e$. For $\epsilon = 3/e - 1$, we get $\frac{1}{n}\sum_{i\in\mathcal{A}'} \frac{V_i(\mathbf{y}_i^*)}{V_i(\mathbf{y}_i^{(s)})} \leq 3(2+\frac{4}{c})$ with constant probability. Then, we apply Lemma 6 with $\mathbf{y}' = \mathbf{y}^{(s)}$ and $\beta = 3(2+\frac{4}{c})$. We conclude that there is a matching π such that $\mathrm{NSW}(\mathbf{y}, \pi) \geq \frac{1}{7+12/c} OPT$. ■

Matching recombination: Now that we know a good matching exists, we want to show that there exists another matching ρ with some additional desirable properties. The matching ρ should be such that each agent a either gets significant value from the matching item $\rho(a)$ alone, or there is no item of very large value contributing to agent a in the fractional solution. The solution is a procedure we borrow almost verbatim from [15]: a careful combination of the initial matching τ and a matching π optimal with respect to our fractional solution $\mathbf{y}$. Our goal is to prove the following lemma, analogous to Lemma 6.1 in [15]. Since our setup here is somewhat different, we repeat the whole argument in a self-contained manner. Also, we remark that while this is an actual algorithmic step in [15], we only need this procedure in the analysis.

Lemma 8. *Let $d \geq 2$. Let $\tau : \mathcal{A} \to \mathcal{G}$ be the matching maximizing $\prod_{a\in\mathcal{A}} v_i(\tau(a))$, $\mathcal{H} = \tau(\mathcal{A})$ the items allocated in this matching. Let $\mathbf{y} \in [0,1]^{\mathcal{A}'\times\mathcal{G}'}$ and let $\pi : \mathcal{A} \to \mathcal{H}$ be any matching. Then there is a matching $\rho : \mathcal{A} \to \mathcal{H}$ such that*

$$NSW(\mathbf{y}, \rho) \geq \frac{1}{d+2} NSW(\mathbf{y}, \pi)$$

and for every agent $a \in \mathcal{A}$,

- *(i) either $v_a(\rho(a)) \geq \frac{1}{d} V_a(\mathbf{y}_a)$ (in which case the ρ-matching item itself recovers a constant fraction of agent a's value)*
- *(ii) or for every item $j \in \mathcal{G}'$, $v_a(j) < \frac{1}{d} V_a(\mathbf{y}_a)$ (in which case there are no items with large contributions to $V_a(\mathbf{y}_a)$).*

Proof: Let τ be the initial optimal matching and $\mathcal{H} = \tau(\mathcal{A})$. Let $\mathbf{y} \in [0,1]^{\mathcal{A}'\times\mathcal{G}'}$ and let $\pi : \mathcal{A} \to \mathcal{H}$ be any matching. (We will use the optimal matching with respect to $\mathbf{y}$ but that is not relevant now.)

We will construct a new matching ρ which combines τ and π in a certain way. First, whenever $\tau(a) = \pi(a)$, we set $\rho(a) = \tau(a) = \pi(a)$. Next, we consider the two matchings as sets of edges $(a, \pi(a))$ and $(a, \tau(a))$ and consider their symmetric difference, $\pi\Delta\tau$. The symmetric difference consists of alternating paths and cycles covering the agents such that $\pi(a) \neq \tau(a)$.

Let $\mathcal{B} = \{a \in \mathcal{A} : v_a(\pi(a)) < \frac{1}{d-1} V_a(\mathbf{y}_a)\}$. We define a modified matching π' where $\pi'(a) = \pi(a)$ for $a \notin \mathcal{B}$ and $\pi'(a) = \emptyset$ for $a \in \mathcal{B}$, meaning that agents $a \in \mathcal{B}$ don't get any items in π'. If $\pi(a)$ contributes less than $\frac{1}{d-1} V_a(\mathbf{y}_a)$, we have

$$\begin{aligned} V_a(\mathbf{y}_a + \mathbf{1}_{\pi(a)}) &\leq V_a(\mathbf{y}_a) + v_a(\pi(a)) \\ &\leq \frac{d}{d-1} V_a(\mathbf{y}_a) \leq \frac{d}{d-1} V_a(\mathbf{y}_a + \mathbf{1}_{\pi'(a)}), \end{aligned}$$

and so

$$\begin{aligned} \mathrm{NSW}(\mathbf{y}, \pi') &= \left(\prod_{a\in\mathcal{A}} V_a(\mathbf{y}_a + \mathbf{1}_{\pi'(a)})\right)^{1/n} \\ &\geq \left(\prod_{a\in\mathcal{A}} \frac{d-1}{d} V_a(\mathbf{y}_a + \mathbf{1}_{\pi(a)})\right)^{1/n} \\ &= \frac{d-1}{d} \mathrm{NSW}(\mathbf{y}, \pi). \end{aligned}$$

Consider an alternating path/cycle C in $\pi\Delta\tau$ and its set of agents $\mathcal{A}(C)$. We distinguish two cases.

1) $\mathcal{B} \cap \mathcal{A}(C) = \emptyset$ (π provides good value for all agents in $\mathcal{A}(C)$). In this case we set $\rho(a) = \pi(a)$ for all $a \in \mathcal{A}(C)$.
2) $\mathcal{B} \cap \mathcal{A}(C) \neq \emptyset$ (some agents in $\mathcal{A}(C)$ don't get good value from π). We remove from C every edge $(a, \pi(a))$ such that $a \in \mathcal{B}$ (which means that $\pi'(a) = \emptyset$); this breaks C into alternating paths. Let us consider one such alternating path, denoting the agents on it $a_1, a_2, \ldots, a_k$ and the items $i_1, i_2, \ldots, i_k$. If $k = 1$, the path consists of just one edge (a_1, i_1). If $k > 1$, the path consists of edges $(a_1, i_1), (i_1, a_2), (a_2, i_2), \ldots, (a_k, i_k)$, where $i_j = \tau(a_j)$ for $j \leq k$ and $i_j = \pi(a_{j+1})$ for $j < k$. We also have $a_1 \in \mathcal{B}$ (this is an agent who does not get any item in π') and $a_2, \ldots, a_k \notin \mathcal{B}$.
   We use the following criterion to decide whether we should use the π-edges or the τ-edges from this alternating path: Let

$$\begin{aligned} &\varphi(a_1, \ldots, a_k) \\ =&\frac{V_{a_1}(\mathbf{y}_{a_1})}{V_{a_1}(\mathbf{y}_{a_1} + \mathbf{1}_{\tau(a_1)})} \prod_{j=2}^{k} \frac{V_{a_j}(\mathbf{y}_{a_j} + \mathbf{1}_{\pi(a_j)})}{V_{a_j}(\mathbf{y}_{a_j} + \mathbf{1}_{\tau(a_j)})} \\ =&\frac{V_{a_1}(\mathbf{y}_{a_1})}{V_{a_1}(\mathbf{y}_{a_1} + \mathbf{1}_{i_1})} \prod_{j=2}^{k} \frac{V_{a_j}(\mathbf{y}_{a_j} + \mathbf{1}_{i_{j-1}})}{V_{a_j}(\mathbf{y}_{a_j} + \mathbf{1}_{i_j})}. \end{aligned}$$

   This is the factor incurred in the objective function if we switch from π' to τ on this alternating path. We call this alternating path *τ-favorable*[1], if $\varphi(a_1, \ldots, a_k) \leq d^k$, and we define $\rho(a_j) = \tau(a_j) = i_j$ for $1 \leq j \leq k$. Otherwise, we call it π-favorable and we define $\rho(a_1) = \emptyset$, $\rho(a_j) = \pi(a_j) = i_{j-1}$ for $2 \leq j \leq k$.
   If we view the process as starting from the matching π' and then applying a swap for each τ-favorable path,

[1] "reversible" in [15]

we obtain a solution $(\mathbf{y}, \rho)$ of value

$$\begin{aligned}&\mathrm{NSW}(\mathbf{y}, \rho)\\ &= \left(\prod_{(a_1,\ldots,a_k)\in\mathcal{P}_\tau} \varphi(a_1,\ldots,a_k)\right)^{-\frac{1}{n}} \mathrm{NSW}(\mathbf{y}, \pi')\\ &\geq \left(\prod_{(a_1,\ldots,a_k)\in\mathcal{P}_\tau} d^k\right)^{-\frac{1}{n}} \mathrm{NSW}(\mathbf{y}, \pi')\end{aligned}$$

where $\mathcal{P}_\tau$ is the set of τ-favorable alternating paths. Since the alternating paths are disjoint in terms of the agents they cover, $\prod_{(a_1,\ldots,a_k)\in\mathcal{P}_\tau} d^k \leq d^n$, and we obtain that

$$\begin{aligned}\mathrm{NSW}(\mathbf{y}, \rho) &\geq \frac{1}{d}\mathrm{NSW}(\mathbf{y}, \pi')\\ &\geq \frac{d-1}{d^2}\mathrm{NSW}(\mathbf{y}, \pi)\\ &\geq \frac{1}{d+2}\mathrm{NSW}(\mathbf{y}, \pi).\end{aligned}$$

Now we turn to the guarantee for each agent $a \in \mathcal{A}$. If $\rho(a) = \tau(a)$ (i.e. the agent receives an item from the initial matching), then we have either $v_a(\rho(a)) \geq \frac{1}{d}V_a(\mathbf{y}_a)$ which satisfies (i), or by the optimality of the initial matching, we have for every $j \in \mathcal{G} \setminus \mathcal{H}$, $v_a(j) \leq v_a(\tau(a)) = v_a(\rho(a)) < \frac{1}{d}V_a(\mathbf{y}_a)$ which satisfies (ii).

If $\rho(a) = \pi(a)$, then this means that a was on a π-favorable alternating path, and also $a \notin \mathcal{B}$ because otherwise we would have set $\rho(a) = \pi'(a) = \emptyset$. So this means that $v_a(\rho(a)) = v_a(\pi(a)) \geq \frac{1}{d-1}V_a(\mathbf{y}_a)$ by the definition of $\mathcal{B}$. So we satisfy (i).

The last case is that $\rho(a) = \emptyset$. This means that $a = a_1 \in \mathcal{B}$ is the starting point of a π-favorable path P, and $\rho(a) = \pi'(a) = \emptyset$. Consider any item $j \in \mathcal{G}\setminus\mathcal{H}$. In the initial matching τ, we could replace the τ-edges on P by the π-edges, and in addition assign j to agent a_1. However, this would not result in an improvement since τ was optimal (as a stand-alone matching). Therefore, we have the following inequality:

$$\frac{v_{a_1}(j)}{v_{a_1}(\tau(a_1))} \cdot \prod_{j=2}^{k} \frac{v_{a_j}(\pi(a_j))}{v_{a_j}(\tau(a_j))} \leq 1.$$

Recall that $a_2, \ldots, a_k \notin \mathcal{B}$ and therefore $v_{a_j}(\pi(a_j)) \geq \frac{1}{d-1}V_{a_j}(\mathbf{y}_{a_j})$ for $j = 2,\ldots,k$. This implies that $\frac{v_{a_j}(\pi(a_j))}{v_{a_j}(\tau(a_j))} \geq \frac{1}{d}\frac{V_{a_j}(\mathbf{y}_{a_j}+\mathbf{1}_{\pi(a_j)})}{v_{a_j}(\tau(a_j))} \geq \frac{1}{d}\frac{V_{a_j}(\mathbf{y}_{a_j}+\mathbf{1}_{\pi(a_j)})}{V_{a_j}(\mathbf{y}_{a_j}+\mathbf{1}_{\tau(a_j)})}$, and also obviously $\frac{v_{a_1}(j)}{v_{a_1}(\tau(a_1))} \geq \frac{v_{a_1}(j)}{V_{a_1}(\mathbf{y}_{a_1}+\mathbf{1}_{\tau(a_1)})}$. Therefore, we have

$$\frac{v_{a_1}(j)}{V_{a_1}(\mathbf{y}_{a_1}+\mathbf{1}_{\tau(a_1)})} \cdot \prod_{j=2}^{k} \frac{V_{a_j}(\mathbf{y}_{a_j}+\mathbf{1}_{\pi(a_j)})}{V_{a_j}(\mathbf{y}_{a_j}+\mathbf{1}_{\tau(a_j)})} \leq d^{k-1}.$$

Finally, since the path is π-favorable, we have

$$\begin{aligned}&\varphi(a_1, i_1, \ldots, a_k, i_k)\\ &= \frac{V_{a_1}(\mathbf{y}_{a_1})}{V_{a_1}(\mathbf{y}_{a_1}+\mathbf{1}_{\tau(a_1)})} \prod_{j=2}^{k} \frac{V_{a_j}(\mathbf{y}_{a_j}+\mathbf{1}_{\pi(a_j)})}{V_{a_j}(\mathbf{y}_{a_j}+\mathbf{1}_{\tau(a_j)})} > d^k.\end{aligned}$$

Combining the last two inequalities, we obtain

$$v_{a_1}(j) < \frac{1}{d}V_{a_1}(\mathbf{y}_{a_1})$$

which means that agent $a = a_1$ satisfies (ii). ∎

F. Conclusion of the analysis

We conclude the analysis by showing that the matching ρ we proved to exist in Section III-E provides a good value with our fractional solution, even if we ignore the contribution of large items. Hence we can obtain an integral assignment which provides a constant-factor approximation relative to OPT and thus prove Theorem 1.

Lemma 9. *Let $(S_1, \ldots, S_n)$ be the assignment obtained by* $\mathbf{RestrictedRandomizedRounding}(\mathbf{y})$ *with parameter $c > 0$ and $\mathbf{y}_i^{(s)} = \mathbf{y}_i^{(L_i)} + \mathbf{1}_{S_i}$ the sparsified fractional solution. Then there exists a matching $\rho : \mathcal{A} \to \mathcal{H}$ such that*

$$\left(\prod_{i\in\mathcal{A}} v_i(S_i + \rho(i))\right)^{1/n} \geq \frac{OPT}{(7+12/c)(c+3)(c+4)}.$$

Proof: Given the sparsified solution $\mathbf{y}^{(s)}$ and the matching π provided by Corollary 7, satisfying

$$\mathrm{NSW}(\mathbf{y}^{(s)}, \pi) \geq \frac{1}{7+12/c}OPT,$$

let ρ be the matching provided by Lemma 8 with parameter $d = c + 2$. This matching satisfies

$$\begin{aligned}\mathrm{NSW}(\mathbf{y}^{(s)}, \rho) &\geq \frac{1}{c+4}\mathrm{NSW}(\mathbf{y}^{(s)}, \pi)\\ &\geq \frac{1}{(7+12/c)(c+4)}OPT\end{aligned}$$

and for every agent $i \in \mathcal{A}$, either (i) $v_i(\rho(i)) \geq \frac{1}{c+2}V_i(\mathbf{y}_i^{(s)})$ or (ii) for every item $j \in \mathcal{G}' = \mathcal{G}\setminus\mathcal{H}$, $v_i(j) < \frac{1}{c+2}V_i(\mathbf{y}_i^{(s)})$.

For every agent $i \in \mathcal{A}$, if (i) is the case, then we know that

$$\begin{aligned}v_i(S_i + \rho(i)) &\geq v_i(\rho(i))\\ &\geq \frac{1}{c+3}\left(V_i(\mathbf{y}_i^{(s)}) + v_i(\rho(i))\right)\\ &\geq \frac{V_i(\mathbf{y}_i^{(s)} + \mathbf{1}_{\rho(i)})}{c+3}.\end{aligned}$$

Otherwise in case (ii), we have

$$\begin{aligned} v_i(S_i + \rho(i)) &\geq V_i(\mathbf{y}_i^{(s)} + \mathbf{1}_{\rho(i)}) - \sum_{j \in L_i} y_{ij} v_i(j) \\ &\geq V_i(\mathbf{y}_i^{(s)} + \mathbf{1}_{\rho(i)}) - (c+1)\frac{V_i(\mathbf{y}_i^{(s)})}{c+2} \\ &\geq \frac{V_i(\mathbf{y}_i^{(s)} + \mathbf{1}_{\rho(i)})}{c+2}, \end{aligned}$$

where the second inequality holds because $\sum_{j \in L_i} y_{ij} \leq c+1$. In conclusion, we know that for any agent $i \in \mathcal{A}$,

$$v_i(S_i + \rho(i)) \geq \frac{1}{c+3} V_i(\mathbf{y}_i^{(s)} + \mathbf{1}_{\rho(i)}).$$

Therefore, we obtain

$$\begin{aligned} \left(\prod_{i \in \mathcal{A}} v_i(S_i + \rho(i))\right)^{1/n} &\geq \frac{1}{c+3} \mathrm{NSW}(\mathbf{y}^{(s)}, \rho) \\ &\geq \frac{OPT}{(7 + 12/c)(c+3)(c+4)}. \end{aligned}$$

■

Now we can prove the main theorem.

Proof of Theorem 1:

By Lemma 9, there is a matching $\rho : \mathcal{A} \to \mathcal{H}$ such that even if we count only the contribution of the small items S_i allocated in $\mathbf{RestrictedRandomizedRounding}(\mathbf{y})$, we have

$$\left(\prod_{i \in \mathcal{A}} v_i(S_i + \rho(i))\right)^{1/n} \geq \frac{OPT}{(7 + 12/c)(c+3)(c+4)}.$$

This means that the same holds for the sets R_i allocated in our algorithm by $\mathbf{RandomizedRounding}(\mathbf{y})$, since $S_i \subseteq R_i$ (the sets R_i include additionally the large items after rounding). In the final step, we find a matching σ which is at least as good as ρ. We choose $c = 1$ which gives

$$\begin{aligned} \left(\prod_{i \in \mathcal{A}} v_i(R_i + \sigma(i))\right)^{1/n} &\geq \left(\prod_{i \in \mathcal{A}} v_i(S_i + \rho(i))\right)^{1/n} \\ &\geq \frac{OPT}{19 \cdot 4 \cdot 5} = \frac{OPT}{380}. \end{aligned}$$

■

IV. Conclusion

We have shown a constant-factor approximation algorithm for Nash Social Welfare with submodular valuations, which is the largest natural class of valuations that allows a constant-factor approximation (using value queries) even for additive welfare maximization. However, there are still several directions and open problems to explore. An obvious one is to improve the approximation ratio which is rather large. As we mentioned, we believe that a substantially smaller (say double-digit) factor is hard to achieve with our approach.

Another open problem is the asymmetric Nash Social Welfare problem, where the objective function is a weighted geometric mean of the agents' valuation functions: $\prod_{i=1}^{n} (v_i(S_i))^{\omega_i}$ for some $\omega_i \geq 0$ (the problem we consider is $\omega_i = 1/n$). The goal is to get a constant-factor approximation independent of the weights ω_i. We remark that [15] gives an approximation guarantee dependent on the weights ω_i; we do not pursue this direction here. For the asymmetric problem, getting a universal constant factor is open even in the the basic case of additive valuations.

Last but not least, solutions optimizing Nash Social Welfare often have additional fairness properties like the envy-free property, or envy-freeness up to one good (see [1]). A line of work has been developed in trying to achieve approximation guarantees for Nash Social Welfare and certain fairness guarantees at the same time [5], [8], [9]. However, our solution does not seem to have such properties and constant-factor approximations with additional fairness guarantees are still unknown for valuation classes beyond additive ones.

References

[1] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. Approximation algorithms for computing maximin share allocations. *ACM Trans. Algorithms*, 13(4):52:1–52:28, 2017.

[2] Nima Anari, Tung Mai, Shayan Oveis Gharan, and Vijay V. Vazirani. Nash social welfare for indivisible items under separable, piecewise-linear concave utilities. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pages 2274–2290, 2018.

[3] Nima Anari, Shayan Oveis Gharan, Amin Saberi, and Mohit Singh. Nash social welfare, matrix permanent, and stable polynomials. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

[4] Siddharth Barman, Umang Bhaskar, Anand Krishna, and Ranjani G. Sundaram. Tight approximation algorithms for p-mean welfare under subadditive valuations. In *28th Annual European Symposium on Algorithms, ESA 2020*, volume 173 of *LIPIcs*, pages 11:1–11:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

[5] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 557–574, 2018.

[6] Gruia Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.

[7] Bhaskar Ray Chaudhury, Yun Kuen Cheung, Jugal Garg, Naveen Garg, Martin Hoefer, and Kurt Mehlhorn. On fair division for indivisible items. In *38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

[8] Bhaskar Ray Chaudhury, Jugal Garg, and Ruta Mehta. Fair and efficient allocations under subadditive valuations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5269–5276, 2021.

[9] Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. A little charity guarantees almost envy-freeness. *SIAM Journal on Computing*, 50(4):1336–1358, 2021.

[10] Richard Cole, Nikhil Devanur, Vasilis Gkatzelis, Kamal Jain, Tung Mai, Vijay V Vazirani, and Sadra Yazdanbod. Convex program duality, Fisher markets, and Nash social welfare. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 459–460, 2017.

[11] Richard Cole and Vasilis Gkatzelis. Approximating the Nash social welfare with indivisible items. *SIAM Journal on Computing*, 47(3):1211–1236, 2018.

[12] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. *Math. Oper. Res.*, 35(1):1–13, 2010.

[13] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey. An analysis of approximations for maximizing submodular set functions – II. *Mathematical Programming Study*, 8:73–87, 1978.

[14] Jugal Garg, Martin Hoefer, and Kurt Mehlhorn. Approximating the Nash social welfare with budget-additive valuations. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2326–2340. SIAM, 2018.

[15] Jugal Garg, Edin Husić, and László A. Végh. Approximating Nash social welfare under Rado valuations. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 1412–1425. ACM, 2021.

[16] Jugal Garg, Pooja Kulkarni, and Rucha Kulkarni. Approximating Nash social welfare under submodular valuations through (un) matchings. In *Proceedings of the fourteenth annual ACM-SIAM symposium on discrete algorithms*, pages 2673–2687. SIAM, 2020.

[17] Frank Kelly. Charging and rate control for elastic traffic. *Eur. Trans. Telecommun.*, 8(1):33–37, 1997.

[18] Euiwoong Lee. APX-hardness of maximizing Nash social welfare with indivisible items. *Information Processing Letters*, 122:17–20, 2017.

[19] Wenzheng Li and Jan Vondrák. Estimating the Nash social welfare for coverage and other submodular valuations. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1119–1130. SIAM, 2021.

[20] J. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, April 1950.

[21] H.R. Varian. Equality, envy, efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974.

[22] Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the Annual ACM Symposium on Theory of Computing*, pages 67–74, 2008.