A Note on Approximating Weighted Nash Social Welfare with Additive Valuations

Yuda Feng
Department of Computer Science and Technology, Harbin Institute of Technology, Heilongjiang, China

Shi Li
Department of Computer Science and Technology, Nanjing University, Jiangsu, China

Abstract

We give the first $O(1)$-approximation for the weighted Nash Social Welfare problem with additive valuations. The approximation ratio we obtain is $e^{1/e} + \epsilon \approx 1.445 + \epsilon$, which matches the best known approximation ratio for the unweighted case [3].

Both our algorithm and analysis are simple. We solve a natural configuration LP for the problem, and obtain the allocation of items to agents using a randomized version of the Shmoys-Tardos rounding algorithm developed for unrelated machine scheduling problems [30]. In the analysis, we show that the approximation ratio of the algorithm is at most the worst gap between the Nash social welfare of the optimum allocation and that of an EF1 allocation, for an unweighted Nash Social Welfare instance with identical additive valuations. This was shown to be at most $e^{1/e} \approx 1.445$ by Barman et al. [3], leading to our approximation ratio.

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis

Keywords and phrases Nash Social Welfare, Configuration LP, Approximation Algorithms

Digital Object Identifier 10.4230/LIPIcs.ICALP.2024.63

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: https://arxiv.org/abs/2404.15607

Funding *Shi Li*: The work of SL was supported by the State Key Laboratory for Novel Software Technology, and the New Cornerstone Science Laboratory.

Acknowledgements YF is an incoming PhD student at Nanjing University, and his work was a part of his undergraduate dissertation supervised by SL.

1 Introduction

In the weighted (or asymmetric) Nash Social Welfare problem with additive valuations, we are given a set $\mathcal{A}$ of n agents, and a set $\mathcal{G}$ of m indivisible items. Every agent $i \in \mathcal{A}$ has a weight $w_i \geq 0$ such that $\sum_{i\in\mathcal{A}} w_i = 1$. There is a value $v_{ij} \in \mathbb{R}_{\geq 0}$ for every $i \in \mathcal{A}$ and $j \in \mathcal{G}$. The goal of the problem is to find an allocation $\sigma : \mathcal{G} \to \mathcal{A}$ of items to agents so as to maximize the following weighted Nash social welfare of σ:

$$\prod_{i\in\mathcal{A}} \left(\sum_{j\in\sigma^{-1}(i)} v_{ij} \right)^{w_i}.$$

In the case where all w_i's are equal to $\frac{1}{n}$, we call the problem the unweighted (or symmetric) Nash Social Welfare problem.

Allocating resources in a fair and efficient manner among multiple agents is a fundamental problem in computer science, game theory, and economics, with applications across diverse domains [19, 33, 4, 28, 25, 2, 29, 5]. The weighted Nash social welfare function is a notable

51st International Colloquium on Automata, Languages, and Programming (ICALP 2024).
Editors: Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson;
Article No. 63; pp. 63:1–63:9
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

objective that balances efficiency and fairness. The unweighted (or symmetric) objective was independently proposed by different communities [26, 20, 32], and later the study has been extended to the weighted case [16, 18]. Since then it has been used in a wide range of applications, including bargaining theory [21, 7, 31], water allocation [17, 10], and climate agreements [34].

The unweighted Nash Social Welfare problem with additive valuations is proved to be NP-hard by Nguyen et al. [27], and APX-hard by Lee [22]. Later the hardness of approximation was improved to $\sqrt{8/7} \approx 1.069$ by Garg et al. [12], via a reduction from Max-E3-Lin-2.

On the positive side, Cole and Gkatzelis [9] gave a $(2e^{1/e} + \epsilon \approx 2.889 + \epsilon)$-approximation using a market equilibrium with some spending restrictions. The ratio was improved by Cole et al. [8] to 2 using a tight analysis, and by Anari et al. [1] to e via a connection of the problem to real stable polynomials. Both papers formulated some convex program (CP) relaxations for the problem. In particular, [8] showed that the optimum solution to their CP corresponds to the spending-restricted market equilibrium defined in [9]. The state-of-the-art result for the problem is a combinatorial $(e^{1/e} + \epsilon \approx 1.45 + \epsilon)$-approximation algorithm due to Barman et al.[3]. They showed that when all the valuations of agents are identical, any allocation that is envy-free up to one item (EF1) is $e^{1/e}$-approximate. Their approximation result then follows from a connection between the non-identical and identical valuation settings they established.

All the results discussed above are for the unweighted case. For the weighted case with agent weights $w \in [0,1]^{\mathcal{A}}, |w|_1 = 1$, Brown et al. [6] presented a $5 \cdot \exp(2 \cdot D_{\mathrm{KL}}\left(w || \frac{\vec{1}}{n})\right) = 5 \cdot \exp(2 \log n + 2 \sum_{i \in \mathcal{A}} w_i \log w_i)$ approximation algorithm, where D_{KL} denotes the KL divergence of two distributions. This is the first work that studies the weighted version for the additive valuation case. Prior to this work, there is an $O(nw_{\max}) = O(n \max_{i \in \mathcal{A}} w_i)$-approximation for the more general submodular valuation case [13], which we discuss soon. Brown et al. [6] showed that the two CPs from [8] and [1] are equivalent, and their result is based on the CP from [8], generalized to the weighted setting.

The additive valuation setting is a special case of the submodular valuation setting, which is another important setting studied in the literature. In this setting, instead of a v_{ij} value for every ij pair, we are given a monotone submodular function $v_i : 2^{\mathcal{G}} \to \mathbb{R}_{\geq 0}$ for every agent $i \in \mathcal{A}$. Our goal is to find an allocation $\sigma : \mathcal{G} \to \mathcal{A}$ so as to maximize $\prod_{i \in \mathcal{A}} \left(v_i(\sigma^{-1}(i))\right)^{w_i}$. A bulk of the previous work has focused on the unweighted case; that is, $w_i = \frac{1}{n}$ for all $i \in \mathcal{A}$. For this case, Garg et al. [15] proved a hardness of $e/(e-1) \approx 1.5819$ using a reduction from Max-3-Coloring; this is better than the 1.069 hardness for the additive valuation case.

On the positive side, Li and Vondrak [24] extended the techniques of Anari et al. [1], to obtain an $e^3/(e-1)^2$-approximation algorithm for the unweighted Nash Social Welfare problem for a large family of submodular valuations, including coverage functions and linear combinations of matroid rank functions. Later, Garg et al. [14] considered a family of submodular functions called Rado functions, and gave an $O(1)$-approximation for this family using the matching theory and convex program techniques. Li and Vondrak [23] developped the first $O(1)$-approximation for general submodular functions, with an approximation ratio of 380. Recently, Garg et al. [13] presented an elegant 4-approximation local search algorithm for the problem, which is the current best approximation result for the problem. All the results discussed above are for the unweighted case. For the weighted case, Garg et al. [13] gave an $O(nw_{\max})$-approximation, where $n_{\max} = \max_{i \in \mathcal{A}} w_i$. Whether the weighted Nash Social Welfare problem with submodular valuations admits a constant approximation is a big open problem.

Recently, the problem has been studied in an even more general setting, namely, the subadditive valuation setting. Dobzinski et al. [11] gives an $O(1)$-approximation for the unweighted Nash Social Welfare problem in this setting under the demand oracle model.

1.1 Our Result and Techniques

In this note, we give the first $O(1)$-approximation algorithm for the weighted Nash Social Welfare problem with additive valuations:

▶ **Theorem 1.** *For any $\epsilon > 0$, there is a randomized $(e^{1/e} + \epsilon \approx 1.445 + \epsilon)$-approximation algorithm for the weighted Nash Social Welfare problem with additive valuations, with running time polynomial in the size of the input and $\frac{1}{\epsilon}$.*

Our approximation ratio of $e^{1/e} + \epsilon$ matches the best ratio for the unweighted case due to Barman et al. [3]. In contrast, the ratio given by Brown et al. [6] is $5 \cdot \exp(2 \cdot D_{\mathrm{KL}}(w||\frac{\vec{1}}{n}))$, which could be polynomial in n.

Our algorithm is based on a natural configuration LP for the problem, which has not been studied before to the best of our knowledge. The configuration LP contains a $y_{i,S}$ variable for every agent i and subset S of items, indicating if the set of items i gets is S or not. We show that the configuration LP can be solved in polynomial time to any precision, despite having exponential number of variables. Once we obtain the LP solution, we define x_{ij} for every $i \in \mathcal{A}$ and $j \in \mathcal{G}$ to be the fraction of j assigned to i.

We use a randomized version of the Shmoys-Tardos rounding algorithm [30] developed for unrelated machine scheduling problems, to round x into an integral solution. For every agent i, we break the fractional items assigned to i into groups from the most valuable to the least, each containing 1 fractional item. The rounding algorithm maintains marginal probabilities, and the requirement that i gets exactly one item from each group (except for the last one, from which i gets at most one item). In the analysis for each agent i, we construct an instance of the unweighted Nash Social Welfare problem with *identical* additive valuations, that involves many copies of the agent i, along with two alloations $\mathcal{S}$ and $\mathcal{S}'$ to the instance. $\mathcal{S}$ corresponds to the LP solution, and $\mathcal{S}'$ corresponds to the randomized solution given by the rounding algorithm. Thanks to the condition that every group contains one item, the solution $\mathcal{S}'$ is envy-free up to one item (EF1). Using the result of [3] about EF1 allocations, we show that the Nash social welfare of $\mathcal{S}'$ is at least $e^{-1/e}$ times that of $\mathcal{S}$, which eventually leads to our $(e^{1/e} + \epsilon)$-approximation.

We believe the configuration LP could be used in many other settings. We leave as an immediate open problem whether it can give an $O(1)$-approximation for the weighted Nash Social Welfare problem with submodular valuations.

2 $(e^{1/e} + \epsilon)$-Approximation Using Configuration LP

We describe the configuration LP in Section 2.1 and the rounding algorithm in Section 2.2. The analysis is given in Section 2.3.

2.1 The Configuration LP

For convenience, for any value function $v : \mathcal{G} \to \mathbb{R}_{\geq 0}$, we define $v(S) := \sum_{j \in S} v_j$ for every $S \subseteq \mathcal{G}$ to be the total value of items in S according to the value funciton v. In the integer program correspondent to the configuration LP, for every $i \in \mathcal{A}$ and $S \subseteq \mathcal{G}$, we have a variable $y_{i,S} \in \{0, 1\}$ indicating if the set of items assigned to i is S or not. We relax the integer constraint to obtain the following configuration LP:

$$\max \sum_{i\in\mathcal{A},S\subseteq\mathcal{G}} w_i\cdot y_{i,S}\cdot \ln v_i(S) \qquad \text{s.t.} \tag{Conf-LP}$$

$$\sum_{i\in\mathcal{A},S\ni j} y_{i,S}\le 1 \qquad \forall j\in\mathcal{G} \tag{1}$$

$$\sum_{S\subseteq\mathcal{G}} y_{i,S}=1 \qquad \forall i\in\mathcal{A} \tag{2}$$

$$y_{i,S}\ge 0 \qquad \forall i\in\mathcal{A}, S\subseteq\mathcal{G} \tag{3}$$

It is convenient for us to consider the natural logarithm of the Nash social welfare function as the objective, which is $\sum_{i\in\mathcal{A}} w_i\cdot\ln v_i(\sigma^{-1}(i))$. This leads to the objective in (Conf-LP). (1) requires that every item j is assigned to at most one agent, and (2) requires that every agent i is assigned one set of items.

The configuration LP has exponential number of variables, but it can be solved within an additive error of $\ln(1+\epsilon)$ for any $\epsilon>0$, in time polynomial in the size of the instance and $\frac{1}{\epsilon}$. We defer the details to Appendix A. Notice that we are considering the logarithm of Nash social welfare, and the typical $(1+\epsilon)$-multiplicative factor becomes an additive error of $\ln(1+\epsilon)$.

2.2 The Rounding Algorithm

From now on, we assume we have obtained a vector y from solving the LP, described using a list of non-zero coordinates; the value of y to (Conf-LP) is at least the optimum value minus $\ln(1+\epsilon)$. We can assume (1) holds with equalities: $\sum_{i\in\mathcal{A},S\ni j} y_{i,S}=1$ for every $j\in\mathcal{G}$. Then we let $x_{ij}=\sum_{S\ni j} y_{i,S}$ for every $i\in\mathcal{A}$ and $j\in\mathcal{G}$. So $\sum_{i\in\mathcal{A}} x_{ij}=1$ for every $j\in\mathcal{G}$.

In this paragraph, we fix an agent i and break the fractional items assigned to i into a set G_i of groups, each containing 1 fractional item. They are created in non-increasing order of values, as in the Shmoys-Tardos algorithm for unrelated machine scheduling problems. That is, the first group contains the 1 fractional most valuable items assigned to i, the second group contains the 1 fractional most valuable items assigned to i after removing the first group, and so on. Formally, we sort the items in $\mathcal{G}$ in non-increasing order of v_{ij} values, breaking ties arbitrarily. Let $p_i=\lceil\sum_{j\in\mathcal{G}} x_{ij}\rceil$. Then we can find vectors $g^1,g^2,\cdots,g^{p_i}\in[0,1]^{\mathcal{G}}$ satisfying the following properties:

(P1) For every $t\in[1,p_i-1]$, we have $|g^t|_1=1$; for $t=p_i$, we have $|g^t|_1=\sum_{j\in\mathcal{G}} x_{ij}-(p_i-1)\in(0,1]$.
(P2) $\sum_{t=1}^{p_i} g^t_j = x_{ij}$ for every $j\in\mathcal{G}$.
(P3) For every $1\le t<t'\le p_i$, and two items j,j' such that j appears before j' in the ordering, it can not happen that $g^t_{j'}>0$ and $g^{t'}_j>0$.

It is easy to see that $g^1,g^2,\cdots,g^{p_i}$ are uniquely decided by the three conditions. We say each g^t is a group. Let $G_i=\{g^1,g^2,\cdots,g^{p_i}\}$ be the set of all groups constructed for this agent i.

Now we take all agents i into consideration and let $G=\uplus_{i\in\mathcal{A}}G_i$ be the set of all groups constructed.[1] The representations of groups give a fractional matching between the groups G and items $\mathcal{G}$: an item j is matched to a group $g\in[0,1]^{\mathcal{G}}$ with a fraction of g_j. Then each item is matched to an extent of 1, and every group g is matched to an extent of $|g|_1$. So

[1] It is possible that two groups from different sets G_i and $G_{i'}$ have the same vector representation. So we treat G as a multiset and we assume we know which set G_i each group $g\in G$ belongs to.

a group is matched to an extent of 1 if it is not the last group for an agent, and at most 1 otherwise. Therefore, we can efficiently output a randomized (partial-)matching between the groups G and items $\mathcal{G}$ so that the marginal probabilities are maintained:

($\star$) For every group $g \in G$ and item $j \in \mathcal{G}$, we have $\Pr[j \text{ is matched to } g] = g_j$.

($\star$) implies that an item $j \in \mathcal{G}$ is matched with probability 1. If a group g has $|g|_1 = 1$, then it is matched with probability 1.

The matching naturally gives us an allocation of items to agents: If an item $j \in \mathcal{G}$ is matched to some group $g \in G_i$, then we assign j to i. By ($\star$) we know that the probability that j is assigned to i is precisely x_{ij}. Let S_i be the set of items assigned to i in the algorithm; notice that it is random. This finishes the description of the randomized rounding algorithm.

2.3 The Analysis

To analyze our rounding algorithm, we first formally define an EF1 allocation.

▶ **Definition 2.** *Given an instance of the unweighted Nash Social Welfare problem with agents $\mathcal{A}$, items $\mathcal{G}$, and identical additive valuation $v : \mathcal{G} \to \mathbb{R}_{\geq 0}$ for all agents, an allocation $\sigma : \mathcal{G} \to \mathcal{A}$ is said to be envy-free up to one item (EF1), if for every two distinct agents i, i' with $\sigma^{-1}(i') \neq \emptyset$, there exists some $j \in \sigma^{-1}(i')$, such that $v(\sigma^{-1}(i') \setminus j) \leq v(\sigma^{-1}(i))$.*

We use the following result from [3]:

▶ **Theorem 3** ([3])**.** *For the unweighted Nash Social Welfare problem with identical additive valuations, any EF1-allocation is an $\mathrm{e}^{1/e}$-approximate solution.*

With the theorem, we prove the following key lemma:

▶ **Lemma 4.** *For every $i \in \mathcal{A}$, we have*

$$\mathbb{E}\left[\ln v_i(S_i)\right] \geq \sum_{S \subseteq \mathcal{G}} y_{i,S} \cdot \ln v_i(S) \; - \; \frac{1}{e}.$$

Proof. Throughout the proof, we fix the agent i. Let $\Delta > 0$ be an integer, so that every $y_{i,S}$ is an integer multiply of $1/\Delta$, and the probability that $S_i = S$ for any S is also an integer multiply of $1/\Delta$.[2] We consider an instance of the unweighted Nash Social Welfare problem with identical additive valuations. In the instance, there are Δ copies of the agent i, and Δx_{ij} copies of every item $j \in \mathcal{G}$; so all the agents are identical. The $y = (y_{i,S})_{S \subseteq \mathcal{G}}$ vector gives us an allocation $\mathcal{S}$ to the instance: For every $S \subseteq \mathcal{G}$, there are exactly $\Delta y_{i,S}$ agents who get a copy of S. Notice that this is a valid solution, as $\sum_S y_{i,S} = 1$ and $\sum_{S \ni j} y_{i,S} = x_{ij}$ for every item j.

The Nash Social Welfare of the allocation $\mathcal{S}$ is

$$\left(\prod_{S \subseteq \mathcal{G}} v_i(S)^{\Delta y_{i,S}}\right)^{1/\Delta} = \prod_{S \subseteq \mathcal{G}} v_i(S)^{y_{i,S}}.$$

The distribution for S_i also corresponds to an allocation $\mathcal{S}'$ of items to agents: For every $S \subseteq \mathcal{G}$, there are $\Delta \cdot \Pr[S_i = S]$ agents who get a copy of S. Again, this is a valid solution as $\sum_S \Pr[S_i = S] = 1$ and $\sum_{S \ni j} \Pr[S_i = S] = \mathbb{E}[S_i \ni j] = x_{ij}$.

[2] We can assume all $y_{i,S}$ values are rational numbers. Under this condition, it is easy to guarantee that the probabilities are rational numbers.

The Nash Social Welfare of the allocation $\mathcal{S}'$ is

$$\left(\prod_{S\subseteq\mathcal{G}} v_i(S)^{\Delta \Pr[S_i=S]}\right)^{1/\Delta} = \prod_{S\subseteq\mathcal{G}} v_i(S)^{\Pr[S_i=S]}.$$

A crucial property for the solution $\mathcal{S}'$ is that it is EF1. Indeed, if $\Pr[S_i = S] > 0$ for some S, then S contains exactly one item from each group in G_i except for the last one, from which S contains at most one item. Also, the items in the groups G_i are sorted by (P3). So if there are two sets S and S' in the support of the distribution for S_i, and we remove the most valuable item from S', then S beats S' item by item.

Therefore, by Theorem 3, we know that the Nash Social Welfare of $\mathcal{S}'$ is at least $e^{-1/e}$ times that of the optimum allocation for the instance, which is at least that of $\mathcal{S}$. That is,

$$\prod_{S\subseteq\mathcal{G}} v_i(S)^{\Pr[S_i=S]} \geq e^{-1/e}\cdot \prod_{S\subseteq\mathcal{G}} v_i(S)^{y_{i,S}}.$$

Taking logarithm on both sides gives the lemma. ◀

Applying the lemma for every $i \in \mathcal{A}$ and using linearity of expectation, we have

$$\mathbb{E}\left[\sum_{i\in\mathcal{A}} w_i \cdot \ln v_i(S_i)\right] \geq \sum_{i\in S, S\subseteq\mathcal{G}} w_i\cdot y_{i,S}\cdot \ln v_i(S) - \frac{1}{e}.$$

We used that $\sum_{i\in\mathcal{A}} w_i = 1$.

By the convexity of exponential function, we have

$$\mathbb{E}\left[\prod_{i\in\mathcal{A}} v_i(S_i)^{w_i}\right] \geq e^{-1/e}\cdot \exp\left(\sum_{i\in S, S\subseteq\mathcal{G}} w_i\cdot y_{i,S}\cdot \ln v_i(S_i)\right) \geq e^{-1/e}\cdot\frac{\text{opt}}{1+\epsilon},$$

where opt is the Nash Social Welfare of the optimum allocation, and the second inequality used that the value of our solution y to (Conf-LP) is at least its optimum value minus $\ln(1+\epsilon)$. By scaling ϵ down by an absolute constant at the beginning, we can make the right side to be at least $\frac{\text{opt}}{e^{1/e}+\epsilon}$. This finishes the proof of Theorem 1.

Finally, we briefly discuss how to derandomize the rounding algorithm. We round the solution to the configuration LP in iterations, maintaining a fractional assignment $\bar{x}$ of items to agents; $\bar{x} = x$ initially. Let Δ be a large enough integer so that every $y_{i,S}$ is an integer multiply of $1/\Delta$. Focus on a fixed agent $i \in \mathcal{A}$ and consider the Nash Social Welfare instance containing Δ copies of i, and $\Delta\bar{x}_{ij}$ copies of each item $j \in \mathcal{G}$. Group the items as follows: the Δ most valuable items belong to the first group, the next Δ most valuable items belong to the second group, and so on. We define Φ_i to be the logarithm of the Nash Social Welfare of the worst allocation satisfying the following condition: every agent gets at most one item from each group. Fortunately, the worst allocation can be defined naturally: the first agent takes the most valuable item from each group, and the second agent takes the second most valuable item from each group, and so on. Thus Φ_i can be computed efficiently. We define $\Phi = \sum_i w_i\Phi_i$ to be the overall potential function. In the randomized version of the algorithm, one can define the rotation operation over the fractional matching between groups G and items $\mathcal{G}$. In expectation the operation does not decrease Φ. To derandomize the algorithm, we can perform the operation deterministically so that Φ does not decrease. The potential value Φ at the end of the algorithm is at least that at the beginning, which is at least the value of the configuration LP minus $1/e$. On the other hand, the logarithm of the Nash Social Welfare of the integral solution is exactly the final Φ. Therefore, the Nash Social Welfare is at least $e^{-1/e}$ times the exponential of the value of the configuration LP.

References

1 Nima Anari, Shayan Oveis Gharan, Amin Saberi, and Mohit Singh. Nash Social Welfare, Matrix Permanent, and Stable Polynomials. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, volume 67 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 36:1–36:12, Dagstuhl, Germany, 2017. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2017.36.

2 Julius Barbanel and Alan Taylor. The Geometry of Efficient Fair Division. *The Geometry of Efficient Fair Division*, pages 1–462, January 2005. doi:10.1017/CBO9780511546679.

3 Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding Fair and Efficient Allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, pages 557–574, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3219166.3219176.

4 Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.

5 Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, USA, 1st edition, 2016.

6 Adam Brown, Aditi Laddha, Madhusudhan Reddy Pittu, and Mohit Singh. Approximation Algorithms for the Weighted Nash Social Welfare via Convex and Non-Convex Programs. In *Proceedings of the Thirty-Fifth ACM-SIAM Symposium on Discrete Algorithms*, 2024.

7 Suchan Chae and Hervé Moulin. Bargaining Among Groups: An Axiomatic Viewpoint. *International Journal of Game Theory*, 39(1):71–88, 2010. doi:10.1007/s00182-009-0157-6.

8 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*, EC '17, pages 459–460, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3033274.3085109.

9 Richard Cole and Vasilis Gkatzelis. Approximating the Nash Social Welfare with Indivisible Items. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '15, pages 371–380, New York, NY, USA, 2015. Association for Computing Machinery. doi:10.1145/2746539.2746589.

10 Dagmawi Mulugeta Degefu, Weijun He, Liang Yuan, An Min, and Qi Zhang. Bankruptcy to Surplus: Sharing Transboundary River Basin's Water under Scarcity. *Water Resources Management: An International Journal, Published for the European Water Resources Association (EWRA)*, 32(8):2735–2751, 2018. URL: https://EconPapers.repec.org/RePEc:spr:waterr:v:32:y:2018:i:8:d:10.1007_s11269-018-1955-z.

11 Shahar Dobzinski, Wenzheng Li, Aviad Rubinstein, and Jan Vondrak. A Constant Factor Approximation for Nash Social Welfare with Subadditive Valuations. *ArXiv*, abs/2309.04656, 2023. arXiv:2309.04656.

12 Jugal Garg, Martin Hoefer, and Kurt Mehlhorn. Satiation in Fisher Markets and Approximation of Nash Social Welfare. *Mathematics of Operations Research*, 0(0):null, 2023. doi:10.1287/moor.2019.0129.

13 Jugal Garg, Edin Husić, Wenzheng Li, László A. Végh, and Jan Vondrák. Approximating Nash Social Welfare by Matching and Local Search. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, pages 1298–1310, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3564246.3585255.

14 Jugal Garg, Edin Husić, and László A. Végh. Approximating Nash Social Welfare under Rado Valuations. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, pages 1412–1425, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3406325.3451031.

15 Jugal Garg, Pooja Kulkarni, and Rucha Kulkarni. Approximating Nash Social Welfare under Submodular Valuations through (Un)Matchings. *ACM Trans. Algorithms*, 19(4), September 2023. doi:10.1145/3613452.

16 John C. Harsanyi and Reinhard Selten. A generalized nash solution for two-person bargaining games with incomplete information. *Management Science*, 18(5):P80–P106, 1972. URL: http://www.jstor.org/stable/2661446.

17 Harold Houba, Gerard van der Laan, and Yuyu Zeng. Asymmetric Nash Solutions in the River Sharing Problem. *Game Theory & Bargaining Theory eJournal*, 2013. URL: https://api.semanticscholar.org/CorpusID:17619205.

18 E. Kalai. Nonsymmetric Nash Solutions and Replications of 2-Person Bargaining. *Int. J. Game Theory*, 6(3):129–133, September 1977. doi:10.1007/BF01774658.

19 Mamoru Kaneko and Kenjiro Nakamura. The Nash Social Welfare Function. *Econometrica*, 47(2):423–435, 1979. URL: http://www.jstor.org/stable/1914191.

20 Frank Kelly. Charging and Rate Control for Elastic Traffic. *European Transactions on Telecommunications*, 8(1):33–37, 1997. doi:10.1002/ett.4460080106.

21 Annick Laruelle and Federico Valenciano. Bargaining in Committees as An Extension of Nash's Bargaining Theory. *Journal of Economic Theory*, 132(1):291–305, 2007. doi:10.1016/j.jet.2005.05.004.

22 Euiwoong Lee. APX-Hardness of Maximizing Nash Social Welfare with Indivisible Items. *ArXiv*, abs/1507.01159, 2015. arXiv:1507.01159.

23 W. Li and J. Vondrak. A Constant-Factor Approximation Algorithm for Nash Social Welfare with Submodular Valuations. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 25–36, Los Alamitos, CA, USA, February 2022. IEEE Computer Society. doi:10.1109/FOCS52979.2021.00012.

24 Wenzheng Li and Jan Vondrák. Estimating the Nash Social Welfare for Coverage and Other Submodular Valuations. In *Proceedings of the Thirty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '21, pages 1119–1130, USA, 2021. Society for Industrial and Applied Mathematics.

25 Herve Moulin. *Fair Division and Collective Welfare.* The MIT Press, 2004.

26 John F. Nash. The Bargaining Problem. *Econometrica*, 18(2):155–162, 1950. URL: http://www.jstor.org/stable/1907266.

27 Nhan-Tam Nguyen, Trung Thanh Nguyen, Magnus Roos, and Jörg Rothe. Complexity and Approximability of Social Welfare Optimization in Multiagent Resource Allocation. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 3*, AAMAS '12, pages 1287–1288, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.

28 Jack Robertson and William Webb. *Cake-Cutting Algorithms: Be Fair if You Can.* A K Peters/CRC Press, 1998.

29 Jrg Rothe. *Economics and Computation: An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division.* Springer Publishing Company, Incorporated, 1st edition, 2015.

30 David B. Shmoys and Éva Tardos. An Approximation Algorithm for the Generalized Assignment Problem. *Mathematical Programming*, 62(1):461–474, 1993. doi:10.1007/BF01585178.

31 William Thomson. Replication Invariance of Bargaining Solutions. *Int. J. Game Theory*, 15(1):59–63, March 1986. doi:10.1007/BF01769276.

32 Hal R Varian. Equity, Envy, and Efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974. doi:10.1016/0022-0531(74)90075-1.

33 H. Peyton Young. *Equity: In Theory and Practice.* Princeton University Press, 1994. URL: http://www.jstor.org/stable/j.ctv10crfx7.

34 S. Yu, E. C. van Ierland, H. P. Weikard, and X. Zhu. Nash Bargaining Solutions for International Climate Agreements under Different Sets of Bargaining Weights. *International Environmental Agreements: Politics, Law and Economics*, 17(5):709–729, 2017. doi:10.1007/s10784-017-9351-3.

A Solving Configuration LP within an Additive Error of $\ln(1+\epsilon)$

Let $\epsilon > 0$ be upper bounded by a sufficiently small constant (we allow ϵ to be a sub-constant). By only allowing every agent to get one item, we can obtain an m-approximation for the our Nash Social Welfare instance. Then, by making $O\left(\frac{\log m}{\epsilon}\right)$ guesses, we can assume we are given a number o such that the value of (Conf-LP) is in $(o, o+\epsilon/3]$.

We consider the dual of (Conf-LP), with the objective replaced by a constraint.

$$\sum_{j\in\mathcal{G}} \alpha_j + \sum_{i\in\mathcal{A}} \beta_i \le o \tag{4}$$

$$\sum_{j\in S} \alpha_j + \beta_i \ge w_i \cdot \ln v_i(S) \qquad \forall i \in \mathcal{A}, S \subseteq \mathcal{G} \tag{5}$$

$$\alpha_j \ge 0 \qquad \forall j \in \mathcal{G} \tag{6}$$

Since (Conf-LP) has value strictly larger than o, the dual LP (4-6) is infeasible. We design an approximate separation oracle for the LP. Given some $\alpha \in \mathbb{R}^{\mathcal{G}}_{\ge 0}$ and $\beta \in \mathbb{R}^{\mathcal{A}}$ that does not satisfy (5), we can find some $i \in \mathcal{A}$ and $S \subseteq \mathcal{G}$ such that

$$\sum_{j\in S} \alpha_j + \beta_i < w_i \ln\big((1+\epsilon/2)v_i(S)\big).$$

The running time of the oracle is polynomial in the input size and $\frac{1}{\epsilon}$. This can be achieved using the standard dynamic programming technique: For a fixed $i \in \mathcal{A}$, to find the S, we guess the item $j^* \in S$ with the largest v_{ij^*}, coarsen the v_{ij} values based on the guess, and run a dynamic programming to find the S.

So, using the ellipsoid method with the approximate separation oracle, we can find polynomially many half spaces of the form $\sum_{j\in S} \alpha_j + \beta_i \ge w_i \ln\big((1+\epsilon/2)v_i(S)\big)$, whose intersection is empty. Then, we consider the Nash Social Welfare instance where all v_{ij} values are scaled up by $1+\epsilon/2$, and (Conf-LP) to the instance. By solving the LP restricted to the variables $y_{i,S}$ correspondent to the half spaces (that is, we let all other variables be 0), we obtain a solution y whose value is at least o w.r.t the scaled instance. So, the value of the solution y to (Conf-LP) w.r.t the original instance is at least $o - \sum_{i\in\mathcal{A}, S\subseteq\mathcal{G}} y_{i,S} w_i \ln(1+\epsilon/2) = o - \sum_i w_i \ln(1+\epsilon/2) = o - \ln(1+\epsilon/2)$.

As the value of (Conf-LP) is at most $o+\epsilon/3$, we solved the LP up to an additive error of $\epsilon/3 + \ln(1+\epsilon/2)$. For a small enough ϵ, this is at most $\ln(1+\epsilon)$.