

Budget-feasible Maximum Nash Social Welfare is Almost Envy-free

Xiaowei Wu¹, Bo Li², Jiarui Gan³

¹IOTSC, University of Macau

²Department of Computing, The Hong Kong Polytechnic University

³Max Planck Institute for Software Systems

xiaoweiwu@um.edu.mo, comp-bo.li@polyu.edu.hk, jrgan@mpi-sws.org

Abstract

The *Nash social welfare* (NSW) is a well-known social welfare measurement that balances individual utilities and the overall efficiency. In the context of fair allocation of indivisible goods, it has been shown by Caragiannis et al. (EC 2016 and TEAC 2019) that an allocation maximizing the NSW is envy-free up to one good (EF1). In this paper, we are interested in the fairness of the NSW in a budget-feasible allocation problem, in which each item has a cost that will be incurred to the agent it is allocated to, and each agent has a budget constraint on the total cost of items she receives. We show that a budget-feasible allocation that maximizes the NSW achieves a 1/4-approximation of EF1 and the approximation ratio is tight. The approximation ratio improves gracefully when the items have small costs compared with the agents’ budgets; it converges to 1/2 when the budget-cost ratio approaches infinity.

1 Introduction

Fairness and efficiency are two of the primary considerations in public economics. Yet, these two objectives are often unaligned and even conflict with each other. The *Nash social welfare* (NSW) — proposed by Nash as a solution for bargaining problems [Nash Jr, 1950; Kaneko and Nakamura, 1979] — is a well-known social welfare function that balances individual utilities and the overall efficiency. In more details, the NSW of a solution \mathbf{x} is defined as the product of the agents’ values $v_i(\mathbf{x})$ provided by this solution. As noted in [Brânzei et al., 2017], it can be viewed as a special case of a family of functions known as *generalized (power) means*:

$$M_p(\mathbf{x}) = \left(\frac{1}{n} \sum_i v_i(\mathbf{x})^p\right)^{1/p}.$$

When $p = 1$, $M_p(\mathbf{x})$ defines the utilitarian social welfare, i.e., the average of all agents’ values; and when $p \rightarrow -\infty$, $M_p(\mathbf{x})$ defines the egalitarian social welfare, i.e., the minimum value received by an agent. The NSW corresponds to exactly the limit of $M_p(\mathbf{x})$ when $p \rightarrow 0$, i.e., $(\prod_{i \in N} v_i(\mathbf{x}))^{1/n}$. Thus, it serves a middle ground between the two other cases. A Max-NSW allocation, i.e., an allocation that maximizes the NSW,

naturally leads to balanced values among the agents while it also nudges the overall efficiency towards the maximum.

In the context of fair allocation of indivisible goods, it has been shown by Caragiannis et al. [Caragiannis et al., 2016; Caragiannis et al., 2019b] that a Max-NSW allocation is *envy-free up to one good* (EF1); namely, it ensures that every agent does not envy the bundle of any other agent by more than one item.

1.1 Our Results

In this paper, we study the fairness of the NSW in a budget-feasible allocation problem. In this problem, every item j has a cost of c_j that will be incurred to the agent that it is allocated to, and every agent i has a budget $B_i > 0$ that can be used to pay for the additive cost of the bundle S of items she receives, i.e., $c(S) = \sum_{j \in S} c_j \leq B_i$. An allocation (X_1, \dots, X_n) is a collection of disjoint subsets of items, whereby X_i is the bundle of items allocated to agent i , yielding value $v_i(X_i)$ for this agent. In our model, we assume that unused budget cannot be converted to an agent’s value. This is true in many real-world applications, e.g., when the cost is the processing time of a job while the value is the payment for finishing the job. With this assumption our model also subsumes the model proposed by Biswas and Barman [Biswas and Barman, 2018], where all items have size 1. An allocation is *budget-feasible* if $c(X_i) \leq B_i$ for all i , i.e., every agent can afford the cost incurred by this allocation. Accordingly, a Max-NSW feasible allocation is the optimal solution to the following problem.

$$\begin{aligned} & \max_{(X_1, \dots, X_n)} \quad \prod_{i \in N} v_i(X_i) \\ & \text{subject to} \quad (X_1, \dots, X_n) \text{ is budget-feasible.} \end{aligned}$$

Note that since allocations are subject to budget constraints, it is not always possible to allocate all items. The unallocated items are assumed to be allocated to a *charity* who has unlimited budget. Indeed, it is not hard to see that a Max-NSW allocation is Pareto-optimal (PO), i.e., we cannot hope to further improve some agent’s utility without hurting the other agents. Here we do not consider the charity as an agent. This immediately implies that no agent has any incentive to exchange her bundle with any subset of unallocated items that fits within her budget. An interesting question is whether some agent would wish to exchange her bundle with other agents.

To answer this question, we adapt the classic EF (envy-free) notion to the budget-feasible setting. In our EF notion, an agent i envies another agent j if there is a subset of items allocated to agent j that costs at most B_i and gives agent i a strictly higher value. The EF1 notion can be defined similarly: agent i envies agent j for more than one item if there is a subset of X_j that costs at most B_i and gives i a strictly higher value even after removing agent i 's most valuable item from this subset. An allocation is said to be EF (resp., EF1) if no agent envies any other agent (resp., for more than one item). We show that, unlike the setting without budget constraints, a Max-NSW allocation is no longer always EF1. However, it remains a good approximation of EF1. Our main contributions are summarized by the following two results, where we call an allocation α -EF1 if it achieves an α approximation of EF1, for $\alpha \in [0, 1]$.

Main Result 1. *A Max-NSW allocation is $\frac{1}{4}$ -EF1 and PO.*

We also show that the approximation ratio we proved cannot be improved. Specifically, we construct an instance for which a Max-NSW allocation is exactly $1/4$ -EF1. We observe that this instance requires some items to have very high costs compared with the agents' budgets, i.e., the budget-cost ratio is $\kappa = \min_{i,j} \frac{B_i}{c_j} = 1$. This motivates our study of the *large budget case* where κ is large. Interestingly, we find that the approximation guarantee improves gracefully with κ and eventually converges to $1/2$ when κ goes to infinity.

Main Result 2. *A Max-NSW allocation is $(\frac{1}{2} - \frac{5}{\kappa^{1/4}})$ -EF1.*

Note that instances with large budget are also a typical setting of the bin packing problem that has received considerable interest in the literature [Feldman *et al.*, 2010; Devanur *et al.*, 2011; Molinaro and Ravi, 2012; Kesselheim *et al.*, 2014]. We remark that when κ goes to infinity our model does not degenerate to the one without budget constraints, because the number of items can increase to infinity as well. This is why the approximation ratio does not converge to 1 as in the setting without budget constraints.

Technical Novelty. We remark that deriving the above results requires techniques very different from those employed in previous work. The approach of Caragiannis *et al.* [Caragiannis *et al.*, 2016; Caragiannis *et al.*, 2019b] for proving the EF1-ness of Max-NSW allocations is by contradiction: if a Max-NSW allocation were not EF1, then by reallocating some items, the NSW could be improved. In our setting, such reallocations need to conform to the budget constraints. We may need to discard some items when reallocating items to an agent. Thus, applying the same argument requires a larger difference between the agents' values. Indeed, this is why the Max-NSW allocation is only approximately EF1 in our setting, instead of being exactly EF1. Our analysis for the second main result is more challenging. We observe that if the items are divisible, then any Max-NSW allocation must be $1/2$ -EF1. To convert the analysis to the indivisible case, we show that in the divisible analysis at most a constant number of items are fractionally allocated. Then it suffices to give a rounding scheme for a small number of fractional items. A main challenge arises when we have to deal with fractional

items with large values. We resolve this issue by using a fine-tuned classification of items into heavy and light ones and applying different analyses depending on the class the majority of the items belong to.

1.2 Related Work

The pioneering work of Caragiannis *et al.* [Caragiannis *et al.*, 2016; Caragiannis *et al.*, 2019b] triggered a series of follow-up interests in the fairness of the NSW. It is shown in [Biswas and Barman, 2018] that with matroid constraints a Max-NSW allocation is always EF1 if the agents have identical valuations. More recently, Amanatidis *et al.* [Amanatidis *et al.*, 2020] proved that a Max-NSW allocation guarantees envy-freeness up to *any* item (EFX) if there are at most two possible values for the goods. The fairness of NSW has also been studied in the setting when the goods are public [Conitzer *et al.*, 2017; Fain *et al.*, 2018]. Furthermore, the algorithm in [Barman *et al.*, 2018] showed that there is an allocation that is simultaneously $1/1.45$ -approximate Max-NSW, Pareto optimal, and $(1 - \epsilon)$ -approximate EF1. Similarly, the algorithm in [McGlaughlin and Garg, 2020] computes an allocation that is $1/2$ -approximate Max-NSW, proportional up to one item, $1/(2n)$ -approximate maximin share fair, and Pareto optimal and can be found in polynomial time. No previous work has considered the setting with budget constraints to the best of our knowledge.

Our work is conceptually related to a line of work on fair division that involves a charity in the models. In our budget-feasible setting, unallocated items are assumed to be donated to a charity. The concept of charity has also been used to simplify analysis of EFX allocations as was in [Caragiannis *et al.*, 2016; Caragiannis *et al.*, 2019b]. Caragiannis *et al.* [Caragiannis *et al.*, 2019a] also showed that there is an EFX allocation on a subset of items that is a $1/2$ -approximation of Max-NSW over the original set of items. Similarly, Chaudhury *et al.* [Chaudhury *et al.*, 2020] proved that by donating no more than $n - 1$ items, there is an EFX allocation for the remaining items and no agent envies the charity.

The complexity of computing a Max-NSW allocation has been studied for various types of valuation functions [Ramezani and Endriss, 2009; Darmann and Schauer, 2015]. The problem is APX-hard for additive valuation functions [Nguyen *et al.*, 2014; Lee, 2017]. There has been some effort towards designing efficient algorithms to compute approximate Max-NSW allocations [Nguyen and Rothe, 2014; Cole and Gkatzelis, 2015; Cole *et al.*, 2017]. The state-of-the-art approximation ratio of NSW for additive valuations is 1.45 by [Barman *et al.*, 2018].

2 Preliminaries

In our model, a set M of m goods will be allocated to a set N of n agents. Every item $j \in M$ has a cost $c_j \geq 0$ and every agent $i \in N$ has a budget $B_i \geq 0$. Each agent i has an additive valuation function $v_i : 2^M \rightarrow \mathbb{R}^+ \cup \{0\}$ on the subsets of items. Equivalently, we can express the valuation function v_i as a vector $v_i = (v_{i1}, \dots, v_{im})$, where $v_{ij} \in \mathbb{R}^+ \cup \{0\}$ denotes agent i 's value for item j . Upon receiving a subset $S \subseteq M$ of items, an agent i obtains an additive value

$v_i(S) = \sum_{j \in S} v_{ij}$; in addition, the allocation incurs an additive cost $c(S) = \sum_{j \in S} c_j$ to agent i , which needs to be paid using her budget. We assume that the values for the items and the cost of the allocation are not mutually convertible. For example, the cost may represent the size of the items and the budget represents the capacity of her knapsack, which cannot be converted from the items' values.

An allocation $\mathbf{X} = (X_0, X_1, \dots, X_n)$ is a partition of the item set M into $n + 1$ bundles, where $X_i \subseteq M$ is the bundle of items allocated to agent i for each $i \in N$ and X_0 is the set of unallocated items (recall that due to the budget constraints, we cannot guarantee $\bigcup_{i=1}^n X_i = M$). The unallocated items can be thought of as a donation to a special agent 0, i.e., a charity, who has an unbounded budget and value 0 for every item (in the sense that she does not envy any other agents no matter what they get). An allocation \mathbf{X} is *budget-feasible* if

$$c(X_i) \leq B_i, \quad \forall i \in N.$$

In the remainder of this paper, we require every allocation to be feasible. When we write ‘‘an allocation’’ we mean ‘‘a feasible allocation’’ unless otherwise stated. For convenience, for a set X and an element j , we will write $X \cup \{j\}$ as $X + j$ and $X \setminus \{j\}$ as $X - j$ throughout the paper.

Definition 2.1 (Max-NSW). An allocation \mathbf{X}^* is a Max-NSW allocation if it maximizes the Nash social welfare, i.e., $\prod_{i \in N} v_i(X_i^*) \geq \prod_{i \in N} v_i(X_i)$ for any allocation \mathbf{X} . For the case when the maximum NSW is zero, the Max-NSW allocation is the one that (1) maximizes the number of agents $|S|$ that have non-zero values; (2) maximizes the NSW of the induced problem instance on the subset of agents S .

In the budget-feasible setting, an agent cannot hope to take a bundle that costs more than her budget. We adapt the EF notion accordingly and define it as follows.

Definition 2.2 (EF). For any $\alpha \in [0, 1]$, an allocation \mathbf{X} is α -approximate envy-free (α -EF) if for any two agents $i, j \in N$ and any $S \subseteq X_j$ with $c(S) \leq B_i$, it holds that $v_i(X_i) \geq \alpha \cdot v_i(S)$. When $\alpha = 1$, \mathbf{X} is also said to be EF.

This definition is consistent with the standard EF notion when every agent's budget is $c(M)$. Since the items are indivisible, (approximate) EF is hard to achieve. One of the most widely studied relaxed notion is envy-free up to one item.

Definition 2.3 (EF1). For any $\alpha \in [0, 1]$, an allocation \mathbf{X} is α -approximate envy-free up to one item (α -EF1) if for any two agents $i, j \in N$ and any $S \subseteq X_j$ with $c(S) \leq B_i$, there exists $j \in S$ such that $v_i(X_i) \geq \alpha \cdot v_i(S - j)$. When $\alpha = 1$, \mathbf{X} is also said to be EF1.

Trivially, an allocation $\mathbf{X} = (M, \emptyset, \dots, \emptyset)$ (i.e., allocating all items to the charity) is already EF and EF1. Apparently, such an allocation is not efficient for the agents by any reasonable measurement of efficiency. Therefore, we consider Pareto optimality as a criterion for efficiency.

Definition 2.4 (PO). An allocation \mathbf{X} is Pareto optimal (PO) if there exists no allocation \mathbf{X}' such that $v_i(X'_i) \geq v_i(X_i)$ for all $i \in N$ and $v_j(X'_j) > v_j(X_j)$ for some $j \in N$.

Note that in a PO allocation \mathbf{X} no agent i can improve her value by exchanging items with the charity (subject to her budget constraint), i.e., for any $S \subseteq X_i \cup X_0$ such that $c(S) \leq B_i$, it holds that $v_i(X_i) \geq v_i(S)$.

3 Max-NSW Allocation is $\frac{1}{4}$ -EF1 and PO

In this section, we present our first main result, which is stated in the theorem below.

Theorem 3.1. A Max-NSW allocation \mathbf{X}^* is $\frac{1}{4}$ -EF1 and PO.

Before presenting the proof, we remark that the $1/4$ approximation ratio is tight, even when agents have identical valuation and budget. We present an instance for which a Max-NSW allocation is not c -EF1, for any constant $c > 1/4$.

Example 3.2. Let $\epsilon > 0$ be arbitrarily close to 0. Consider the instance with two agents who have the identical valuation v and budget $B = 1$. Let there be $1 + \frac{1}{\epsilon}$ items, where $v_1 = 1 + \epsilon$, $c_1 = 1$, $v_j = 4\epsilon$, and $c_j = \epsilon$ for all $j \in \{2, 3, \dots, 1 + \frac{1}{\epsilon}\}$.

It can be verified that the allocation \mathbf{X}^* with $X_1^* = \{1\}$ and $X_2^* = \{2, \dots, 1 + \frac{1}{\epsilon}\}$ maximizes the NSW with $v(X_1^*) \cdot v(X_2^*) = 4(1 + \epsilon)$. However, \mathbf{X}^* is not c -EF1, for any $c > \frac{1}{4}$, because $v(X_1^*) = \frac{1+\epsilon}{4(1-\epsilon)} \cdot v(X_2^* - j)$ for all $j \in X_2^*$.

j	1	2	...	$1 + \frac{1}{\epsilon}$
v_j	$1 + \epsilon$	4ϵ	...	4ϵ
c_j	1	ϵ	...	ϵ

Next we proceed to the proof of Theorem 3.1.

Proof of Theorem 3.1: We first show that we can assume without loss of generality that the maximum NSW is non-zero. If the maximum NSW is zero, let S be the agents with non-zero values in the Max-NSW allocation \mathbf{X}^* . It is not difficult to verify that (i) any agent not in S does not envy the charity (as otherwise we can increase the size of S); (ii) any agent not in S does not envy any other agent by more than one item (again, otherwise we can increase the size of S); (iii) any agent $i \in S$ does not envy any agent $j \notin S$, because if agent i has non-zero valuation on X_j then the Max-NSW on S can be improved. Hence it remains to argue the (approximate) EF1-ness and PO among agents in $S \cup \{0\}$, which reduces to the case when Max-NSW is not zero¹.

It is trivial to see that \mathbf{X}^* is PO because if there is another allocation that increases one agent's value without decreasing any other agent's value, it must also increase the NSW, which contradicts the maximum of NSW under \mathbf{X}^* . In what follows, we prove that \mathbf{X}^* is $\frac{1}{4}$ -EF1 by contradiction. Suppose that a Max-NSW allocation \mathbf{X}^* is not $\frac{1}{4}$ -EF1, we show that there exists another allocation \mathbf{X}' with a strictly larger NSW, which is a contradiction with \mathbf{X}^* being a Max-NSW allocation.

By assumption, there exist two agents, say 1 and 2, such that agent 1 finds a set $T \subseteq X_2$ such that $c(T) \leq B_1$ and

$$v_1(T - j) > 4 \cdot v_1(X_1), \quad \text{for all } j \in T. \quad (1)$$

For ease of description, we rename the items in T as $\{1, 2, \dots, t\}$ such that $v_{2j} \leq v_{2,j+1}$ for any $1 \leq j < t$. Let $T_1 = \{j \in T - t : j \text{ is odd}\}$ and $T_2 = \{j \in T - t : j \text{ is even}\}$. Note that we have one of the following two inequalities:

$$\begin{aligned} v_2(T_1) &\leq v_2(T_2) \leq v_2(T_1 + t) \quad (\text{if } t \text{ is odd}), \text{ or} \\ v_2(T_2) &\leq v_2(T_1) \leq v_2(T_2 + t) \quad (\text{if } t \text{ is even}). \end{aligned}$$

¹Similar arguments appeared in Caragiannis et al. [Caragiannis et al., 2016; Caragiannis et al., 2019b].

In other words, under the valuation of agent 2, if we assign item t to the bundle with smaller value, then the resulting bundle has value at least that of the other bundle.

Now we can construct a new allocation \mathbf{X}' , where

- $X'_i = X_i$ for all $i \geq 3$,
- $X'_1 = T_l$, where $l \in \arg \max_{\ell \in \{1,2\}} v_1(T_\ell)$,
- $X'_2 = X_2 \setminus T_l$, and
- $X'_0 = X_0 \cup X_1$.

In other words, we first remove all the items originally allocated to agent 1. Then we pick the bundle agent 1 prefers in $\{T_1, T_2\}$, and move items in this bundle from X_2 to X_1 . Note that under the valuation of agent 2, the items agent 2 loses have total value at most $\frac{1}{2} \cdot v_2(T)$. In addition, \mathbf{X}' is a feasible allocation as $c(T_l) \leq c(T) \leq B_1$ and $c(X_2 \setminus T_l) \leq c(X_2) \leq B_2$. By Equation (1), we then have

$$v_1(X'_1) \geq \frac{1}{2} \cdot v_1(T - t) > 2 \cdot v_1(X_1).$$

Moreover,

$$v_2(X'_2) = v_2(X_2) - v_2(T_l) \geq \frac{1}{2} \cdot v_2(X_2).$$

Since $v_i(X'_i) = v_i(X_i)$ for all $i \geq 3$, it follows that

$$\prod_{i \in N} v_i(X'_i) > \prod_{i \in N} v_i(X_i),$$

which contradicts the fact that \mathbf{X}^* is Max-NSW. \blacksquare

Remark. While we assume that the valuation functions are additive, our result holds with *sub-additive* valuation functions as well (see the full version of the paper [Wu *et al.*, 2020]).

4 Improved Ratios with Large Budgets

We have shown that in general, the $1/4$ approximation ratio can not be improved. However, one observation is that in the tight instance we presented (Example 3.2), the budgets are such that $B_1 = B_2 = 1 = \max_{j \in M} c_j$. In other words, an agent's budget may be exhausted by allocating only one item to her. Interestingly, we find that when this is not the case, the approximation ratio can be improved.

4.1 Warm-up Analysis

As a warm up, we show that if it takes at least two item to exhaust the budget of any agent, then the approximation ratio can be slightly improved to $\frac{3}{11} \approx 0.273$.

Lemma 4.1. *Suppose that $B_i \geq 2 \cdot c_j$ for all agent $i \in N$ and item $j \in M$. Then a Max-NSW allocation \mathbf{X}^* is $\frac{3}{11}$ -EF1.*

Proof. Suppose for the sake of contradiction that a Max-NSW allocation \mathbf{X}^* is not $\frac{3}{11}$ -EF1. Without loss of generality, let agents 1 and 2 be the two agents that block \mathbf{X}^* from being $\frac{3}{11}$ -EF1. In other words, there exists $T \subseteq X_2$ with $c(T) \leq B_1$ such that for any $j \in T$, we have

$$v_1(T - j) > \frac{11}{3} \cdot v_1(X_1).$$

We use a similar approach to the proof of Theorem 3.1 and partition T into three subsets T_1, T_2 , and $\{t\}$ such that

$$v_2(T_1) \leq v_2(T_2) \leq v_2(T_1 \cup \{t\}).$$

Case 1: If there exists $l \in \{1, 2\}$ such that $v_1(T_l) > 2 \cdot v_1(X_1)$, we construct a new allocation \mathbf{X}' with $X'_i = X_i$ for all $i \geq 3$, $X'_1 = T_l$, $X'_2 = X_2 \setminus T_l$, and $X'_0 = X_0 \cup X_1$. By this construction, we have that $v_1(X'_1) > 2 \cdot v_1(X_1)$ and $v_2(X'_2) \leq \frac{1}{2} \cdot v_2(X_2)$; thus,

$$\prod_{i \in N} v_i(X'_i) > \prod_{i \in N} v_i(X_i),$$

which contradicts the fact that \mathbf{X}^* is a Max-NSW allocation.

Case 2: $\max\{v_1(T_1), v_1(T_2)\} \leq 2 \cdot v_1(X_1)$. Since

$$v_1(T_1) + v_1(T_2) = v_1(T_1 \cup T_2) > \frac{11}{3} \cdot v_1(X_1),$$

we have $\min\{v_1(T_1), v_1(T_2)\} > \frac{5}{3} \cdot v_1(X_1)$. Pick an arbitrary $l \in \arg \min_{\ell \in \{1,2\}} c(T_\ell)$. Since $c(T_1) + c(T_2) \leq B_1$, we have $c(T_l) \leq \frac{1}{2} \cdot B_1$.

We construct a new allocation in which we allocate all the items in T_l to agent 1, along with a sufficiently valuable portion of items in X_1 . To ensure that the budget constraint is not violated, we need the following result.

Claim 4.1. *There exists $\hat{X}_1 \subseteq X_1$ such that $c(\hat{X}_1) \leq \frac{1}{2} \cdot B_1$ and $v_1(\hat{X}_1) \geq \frac{1}{3} \cdot v_1(X_1)$.*

Proof. If $c(X_1) \leq \frac{1}{2} \cdot B_1$, we can let $\hat{X}_1 = X_1$. Note that under this case, $v_1(\hat{X}_1) = v_1(X_1)$. Now suppose $c(X_1) > \frac{1}{2} \cdot B_1$. We create a set T and let it be an empty set initially. Then we add into T items in X_1 one at a time until $c(T) > \frac{1}{2} \cdot B_1$. Let j be the last item added into T . We have

$$c(T) > \frac{1}{2} \cdot B_1 \text{ and } c(T - j) \leq \frac{1}{2} \cdot B_1.$$

Consider the partition of X_1 into three sets: $T - j$, $\{j\}$ and $X_1 \setminus T$. Note that all three sets have cost at most $\frac{1}{2} \cdot B_1$. Let \hat{X}_1 be the set with maximum value (under v_1) among $T - j$, $\{j\}$ and $X_1 \setminus T$. Obviously, we have $v_1(\hat{X}_1) \geq \frac{1}{3} \cdot v_1(X_1)$. \square

By the above claim, we can now construct a new allocation \mathbf{X}' with $X'_i = X_i$ for all $i \geq 3$, $X'_1 = T_l \cup \hat{X}_1$, $X'_2 = X_2 \setminus T_l$, and $X'_0 = X_0 \cup (X_1 \setminus \hat{X}_1)$. This allocation is feasible since

$$c(X'_1) = c(T_l) + c(\hat{X}_1) \leq \frac{B_1}{2} + \frac{B_1}{2} = B_1.$$

By construction, we have that

$$\begin{aligned} v_1(X'_1) &= v_1(T_l) + v_1(\hat{X}_1) \\ &> \frac{5}{3} \cdot v_1(X_1) + \frac{1}{3} \cdot v_1(X_1) = 2 \cdot v_1(X_1), \end{aligned}$$

and $v_2(X'_2) \geq \frac{1}{2} \cdot v_2(X_2)$. Therefore $\prod_{i \in N} v_i(X'_i) > \prod_{i \in N} v_i(X_i)$, which is a contradiction. \square

As we can see from the analysis, the improvement in the approximation ratio comes mainly from the fact that when reallocating items in T_l from X_2 to X_1 , we are able to keep a constant fraction of the items in X_1 (as in Claim 4.1).

It is natural to expect that if we are able to obtain an even finer dissection of X_1 , we should be able to improve the approximation ratio even further. Indeed, this is the case. We show that when all items are very small compared with the budgets, the approximation ratio approaches $1/2$. While the idea is clear, as we will show in the following section, to accomplish the analysis is a highly non-trivial task.

4.2 Large Budget Case

Let $\kappa = \min_{i \in N, j \in M} \lfloor \frac{B_i}{c_j} \rfloor$. Without loss of generality, assume κ is an integer; otherwise we round κ down to the nearest integer. In words, it takes at least κ items to exhaust the budget of any agent.

Theorem 4.2. *The Max-NSW allocation \mathbf{X}^* is $(\frac{1}{2} - \frac{5}{\kappa^{1/4}})$ -EF1, where $\kappa = \min_{i \in N, j \in M} \lfloor \frac{B_i}{c_j} \rfloor$.*

Observe that when $\kappa \rightarrow \infty$ the approximation ratio (with respect to EF1) approaches $1/2$. Before we present the proof of Theorem 4.2, we first show that the approximation ratio is tight. In particular, we give an example with arbitrarily large κ , for which a Max-NSW allocation is $(\frac{1}{2} + O(\frac{1}{\kappa}))$ -EF1.

Example 4.3. *Let there be two agents $\{1, 2\}$ with $B_1 = B_2 = \kappa$. Let there be a set of 2κ items $M = M_1 \cup M_2$, with cost 1 each, where $M_1 = \{1, 2, \dots, \kappa\}$ and $M_2 = \{\kappa + 1, \kappa + 2, \dots, 2\kappa\}$. Let the valuations be $v_{1j} = 1$ for all $j \in M_1$; $v_{1j} = 2$ for all $j \in M_2$; $v_{2j} = 0$ for all $j \in M_1$; $v_{2j} = 2$ for all $j \in M_2$.*

j	1	...	κ	$\kappa + 1$...	2κ
v_{1j}	1	...	1	2	...	2
v_{2j}	0	...	0	2	...	2

Suppose that $x \leq \kappa$ items in M_2 are allocated to agent 2 in a Max-NSW allocation, then it is optimal to allocate the remaining $\kappa - x$ items in M_2 together with x items in M_1 to agent 1. The NSW is thus given by $v_1(X_1) \cdot v_2(X_2) = (2\kappa - x) \cdot 2x$, which is maximized when $x = \kappa$. That is, $X_1 = M_1$ and $X_2 = M_2$. Since for any $j \in X_2$,

$$v_1(X_2 - j) = 2(\kappa - 1) = \frac{2(\kappa - 1)}{\kappa} \cdot v_1(X_1),$$

the Max-NSW allocation is $\frac{\kappa}{2\kappa - 1}$ -EF1, and the approximation ratio approaches $\frac{1}{2}$ when $\kappa \rightarrow \infty$.

4.3 Proof of Theorem 4.2

Next we prove Theorem 4.2. For ease of notation we let $k = \kappa^{1/4}$. Note that it suffices to consider the case when $k > 20$, as otherwise the theorem follows from Theorem 3.1.

Definition 4.4 (Density). *We define the density of item $j \in M$ under the valuation of agent $i \in N$ as $\rho_{ij} = \frac{v_{ij}}{c_j}$.*

Lemma 4.5. *Given two sets of items X and Y , both having cost at most B , and a valuation function v , there exists a subset of items $Z \subseteq X \cup Y$ with $c(Z) \leq B$ and*

$$v(Z) \geq \left(1 - \frac{c(Y)}{B} - \frac{1}{k^4}\right) \cdot v(X) + v(Y).$$

Proof. We prove by constructing the subset Z satisfying the claimed properties. In particular, we compute a subset $X' \subseteq X$ with $c(X') \leq B - c(Y)$, and

$$v(X') \geq \left(1 - \frac{c(Y)}{B} - \frac{1}{k^4}\right) \cdot v(X).$$

Then the lemma follows immediately by setting $Z = X' \cup Y$.

To ease the analysis, we assume without loss of generality that $c(X) = B$. If $c(X) < B$ then imagine that we include into X many items with 0 value and arbitrarily small cost until $c(X) = B$, which does not increase the value of X .

We construct X' by repeatedly removing an item in X with the smallest density (under valuation function v), until $c(X') \leq B - c(Y)$. Since each item has cost at most $\frac{B}{k^4}$ by definition of k , the total cost of items removed in this procedure is at most $c(Y) + \frac{B}{k^4}$.

Since the items removed are among those with the smallest density, their average density is at most the overall density of items in X ; the total value of items removed is at most

$$\left(c(Y) + \frac{B}{k^4}\right) \cdot \frac{v(X)}{c(X)} = \left(\frac{c(Y)}{B} + \frac{1}{k^4}\right) \cdot v(X),$$

Thus we have $v(Z) = v(X') + v(Y) \geq \left(1 - \frac{c(Y)}{B} - \frac{1}{k^4}\right) \cdot v(X) + v(Y)$, as claimed. \square

As before, we prove by contradiction, and assume the allocation is not $(\frac{1}{2} - \frac{5}{k})$ -EF1, say, between agents 1 and 2. We show that this assumption leads to contradictions.

Let $T \subseteq X_2$ with $c(T) \leq B_1$ such that

$$\begin{aligned} v_1(T - j^*) &> \left(\frac{1}{2} - \frac{5}{k}\right)^{-1} \cdot v_1(X_1) \\ &= \left(2 + \frac{20}{k-10}\right) \cdot v_1(X_1), \end{aligned} \quad (2)$$

where $j^* = \arg \max_{j \in T} v_{2j}$ is the item in T with the maximum value under the valuation of agent 2.

Let $\hat{T} = T - j^*$. For each item $j \in \hat{T}$, we refer to $v_{2j}/v_2(\hat{T})$ as the *contribution* of item j to set \hat{T} under valuation of agent 2. The total contribution of items in \hat{T} is 1. Depending on the contributions of items, we partition \hat{T} into a set T_h of *heavy items* and a set T_l of *light items*:

$$T_h := \left\{j \in \hat{T} : \frac{v_{2j}}{v_2(\hat{T})} \geq \frac{1}{k^3}\right\}, \quad \text{and} \quad T_l := \hat{T} \setminus T_h.$$

We first show that under the valuation of agent 1, the contribution of heavy items must be small.

Claim 4.2. *For all $j \in T_h$, we have $\frac{v_{1j}}{v_1(\hat{T})} < \frac{1}{2} \cdot \frac{v_{2j}}{v_2(\hat{T})}$.*

Proof. Suppose that there exists $j \in T_h$ such that $\frac{v_{1j}}{v_1(\hat{T})} \geq \frac{1}{2} \cdot \frac{v_{2j}}{v_2(\hat{T})}$. We show that there exists an allocation (by reallocating item j) that achieves a strictly larger NSW, which is a contradiction. Let $c = \frac{v_2(\hat{T})}{v_{2j}}$. By definition of heavy item, we have $c \leq k^3$. Consider the allocation obtained by moving item j from \hat{T}_h to X_1 . Since $v_{2j} \leq v_{2j^*}$, we have

$$\begin{aligned} v_2(X'_2) &= \left(1 - \frac{v_{2j}}{v_2(\hat{T})}\right) \cdot v_2(X_2) \\ &\geq \left(1 - \frac{v_{2j}}{v_2(\hat{T}) + v_{2j^*}}\right) \cdot v_2(X_2) \geq \left(1 - \frac{1}{c+1}\right) \cdot v_2(X_2). \end{aligned}$$

Note that including j into X_1 may result in a violation of the budget constraint of agent 1. To resolve this issue, we use Lemma 4.5 with $B = B_1$, $v = v_1$, $X = X_1$ and $Y = \{j\}$, and let $X'_1 = Z$ (as specified in Lemma 4.5). Since item j has cost $c_j \leq \frac{B_1}{k^4}$, we have

$$v_1(X'_1) \geq \left(1 - \frac{2}{k^4}\right) \cdot v_1(X_1) + v_{1j}.$$

Recall that by assumption of the proof we have

$$v_{1j} \geq \frac{1}{2} \cdot v_1(\hat{T}) \cdot \frac{v_{2j}}{v_2(\hat{T})} = \frac{1}{2c} \cdot v_1(\hat{T}).$$

By inequality (2), we have

$$\begin{aligned} v_1(X'_1) &\geq \left(1 - \frac{2}{k^4} + \frac{1}{2c} \cdot \left(2 + \frac{20}{k-10}\right)\right) \cdot v_1(X_1) \\ &= \left(1 + \frac{1}{c} + \frac{10}{c(k-10)} - \frac{2}{k^4}\right) \cdot v_1(X_1) \\ &> \frac{c+1}{c} \cdot v_1(X_1), \end{aligned}$$

where the last inequality holds since $c \leq k^3$.

Hence, we have $v_1(X'_1) \cdot v_2(X'_2) > v_1(X_1) \cdot v_2(X_2)$ and $\prod_{i \in N} v_i(X'_i) > \prod_{i \in N} v_i(X_i)$, which is a contradiction. \square

Let $f = \frac{v_2(T_l)}{v_2(\hat{T})}$ be the total contribution of light items to \hat{T} .

By the above claim we have

$$\frac{v_1(T_l)}{v_1(\hat{T})} = 1 - \frac{v_1(T_h)}{v_1(\hat{T})} > 1 - \frac{1}{2} \cdot \frac{v_2(T_h)}{v_2(\hat{T})} = \frac{1}{2} + \frac{f}{2}.$$

In other words, under the valuation of agent 1, the light items have a larger total contribution to \hat{T} . Next, we show the following lemma (the proof of the lemma is included in the full version [Wu *et al.*, 2020]). Roughly speaking, we can partition the light items into k sets with about the same cost and value.

Lemma 4.6. *We can partition T_l into k sets such that each set Y satisfies*

$$v_2(Y) \leq \left(\frac{f}{k} + \frac{4}{k^3}\right) \cdot v_2(\hat{T}), \quad c(Y) \leq \left(\frac{1}{k} + \frac{4}{k^4}\right) \cdot B_1.$$

We partition T_l into k sets as specified in Lemma 4.6, then pick the one most valuable to agent 1 out of the k sets, and reassign it to agent 1. Let Y be the set chosen. Note that

$$\begin{aligned} v_1(Y) &\geq \frac{1}{k} \cdot v_1(T_l) \geq \frac{1}{k} \left(\frac{1}{2} + \frac{f}{2}\right) \cdot v_1(\hat{T}) \\ &\geq \frac{1}{k} \left(\frac{1}{2} + \frac{f}{2}\right) \left(2 + \frac{20}{k-10}\right) \cdot v_1(X_1). \end{aligned}$$

After losing set Y , the remaining value of agent 2 is

$$v_2(X'_2) = v_2(X_2) - v_2(Y) \geq \left(1 - \frac{f}{k} - \frac{4}{k^3}\right) \cdot v_2(X_2), \quad (3)$$

where the inequality holds by Lemma 4.6.

After acquiring set Y , agent 1 has the set of items $X'_1 = X_1 \cup Y$, which might have a larger cost than B_1 . Similarly, we use Lemma 4.5, with $B = B_1$, $v = v_1$ and $X = X_1$. Let $X'_1 = Z$ be the set specified in Lemma 4.6.

As a result, the bundle agent 1 receives has value

$$\begin{aligned} &v_1(X'_1) \\ &\geq \left(1 - \left(\frac{1}{k} + \frac{5}{k^4}\right) + \frac{1}{k} \left(\frac{1}{2} + \frac{f}{2}\right) \left(2 + \frac{20}{k-10}\right)\right) \cdot v_1(X_1) \\ &= \left(1 + \frac{f}{k} + \frac{10(1+f)}{k(k-10)} - \frac{5}{k^4}\right) \cdot v_1(X_1) \\ &> \left(1 + \frac{f}{k} + \frac{9}{k^2}\right) \cdot v_1(X_1), \end{aligned} \quad (4)$$

where in the last inequality we use $f \geq 0$ and $k \geq 20$.

Combining inequalities (4) and (3), we have

$$\begin{aligned} \frac{v_1(X'_1)v_2(X'_2)}{v_1(X_1)v_2(X_2)} &> \left(1 - \frac{f}{k} - \frac{4}{k^3}\right) \cdot \left(1 + \frac{f}{k} + \frac{9}{k^2}\right) \\ &= \left(1 - \left(\frac{f}{k} + \frac{4}{k^3}\right)^2 + \left(\frac{9}{k^2} - \frac{4}{k^3}\right) \left(1 - \frac{f}{k} - \frac{4}{k^3}\right)\right) \\ &\geq \left(1 - \left(\frac{f}{k} + \frac{4}{k^3}\right)^2 + \frac{5}{k^2}\right) > 1. \end{aligned}$$

It follows that $\prod_{i \in N} v_i(X'_i) > \prod_{i \in N} v_i(X_i)$, which is a contradiction.

5 Discussion and Future Directions

We showed that in the presence of budget constraints a Max-NSW allocation may not be EF1 but achieves a constant approximation of EF1. The tight approximation ratio is $1/4$ and in the case of large budgets, the ratio improves with the budget-cost ratio and converges to $1/2$ when this ratio goes to infinity. There are several directions for future work.

First, our results imply that in a budget-feasible setting, a $1/4$ -EF1 allocation always exists and is compatible with Pareto optimality, but it is not known whether an exact EF1 allocation always exists or not and this appears to be a non-trivial problem. However, we do know that every PO allocation is at most $1/2$ -EF1: the hard instance can be obtained by modifying Example 3.2 and making the value of the first item to be 2. Second, the computation of maximum NSW in our setting is an interesting question which we did not study in this paper. Indeed, an APX-hardness for the complexity of this problem can be readily established since the setting without budget can be seen as a special case of our setting (where every agent has budget $c(M)$). Thus, a natural task is to design efficient approximation algorithms. Theorem 3.1 can be extended to show that an α -approximation of the Max-NSW is also $\frac{\alpha}{4}$ -EF1. Hence an approximation algorithm for Max-NSW would also compute an approximate EF1 allocation. However, in the large budget case, an α -approximation of the Max-NSW is not necessarily $\alpha \cdot \left(\frac{1}{2} - \frac{5}{\kappa^{1/4}}\right)$ -EF1 (we include a hard instance in our full version [Wu *et al.*, 2020]). Finally, while we focus on EF in this paper, it would also be interesting to consider other fairness notions such as proportionality and maximin share fairness and study whether a Max-NSW allocation provides any fairness guarantee under these notions.

Acknowledgments

Xiaowei Wu is funded by the Science and Technology Development Fund, Macau SAR (File no. SKL-IOTSC-2021-2023) and the Start-up Research Grant of University of Macau (File no. SRG2020-00020-IOTSC). Bo Li is supported by The Hong Kong Polytechnic University (Grant no. P0034420). Jiarui Gan was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 945719). The authors thank Edith Elkind, Georgios Birmpas, Warut Suksompong, and Alexandros Voudouris for helpful discussions at the early stage of this work.

References

- [Amanatidis *et al.*, 2020] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A. Voudouris. Maximum Nash welfare and other stories about EFX. In *IJCAI*, pages 24–30. ijcai.org, 2020.
- [Barman *et al.*, 2018] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *EC*, pages 557–574. ACM, 2018.
- [Biswas and Barman, 2018] Arpita Biswas and Siddharth Barman. Fair division under cardinality constraints. In *IJCAI*, pages 91–97. ijcai.org, 2018.
- [Brânzei *et al.*, 2017] Simina Brânzei, Vasilis Gkatzelis, and Ruta Mehta. Nash social welfare approximation for strategic agents. In *EC*, pages 611–628. ACM, 2017.
- [Caragiannis *et al.*, 2016] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. In *EC*, pages 305–322. ACM, 2016.
- [Caragiannis *et al.*, 2019a] Ioannis Caragiannis, Nick Gravin, and Xin Huang. Envy-freeness up to any item with high Nash welfare: The virtue of donating items. In *EC*, pages 527–545. ACM, 2019.
- [Caragiannis *et al.*, 2019b] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. *ACM Trans. Economics and Comput.*, 7(3):12:1–12:32, 2019.
- [Chaudhury *et al.*, 2020] Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. A little charity guarantees almost envy-freeness. In *SODA*, pages 2658–2672. SIAM, 2020.
- [Cole and Gkatzelis, 2015] Richard Cole and Vasilis Gkatzelis. Approximating the Nash social welfare with indivisible items. In *STOC*, pages 371–380. ACM, 2015.
- [Cole *et al.*, 2017] Richard Cole, Nikhil R. Devanur, Vasilis Gkatzelis, Kamal Jain, Tung Mai, Vijay V. Vazirani, and Sadra Yazdanbod. Convex program duality, fisher markets, and Nash social welfare. In *EC*, pages 459–460. ACM, 2017.
- [Conitzer *et al.*, 2017] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In *EC*, pages 629–646. ACM, 2017.
- [Darmann and Schauer, 2015] Andreas Darmann and Joachim Schauer. Maximizing nash product social welfare in allocating indivisible goods. *Eur. J. Oper. Res.*, 247(2):548–559, 2015.
- [Devanur *et al.*, 2011] Nikhil R. Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A. Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *EC*, pages 29–38. ACM, 2011.
- [Fain *et al.*, 2018] Brandon Fain, Kamesh Munagala, and Nisarg Shah. Fair allocation of indivisible public goods. In *EC*, pages 575–592. ACM, 2018.
- [Feldman *et al.*, 2010] Jon Feldman, Monika Henzinger, Nitish Korula, Vahab S. Mirrokni, and Clifford Stein. Online stochastic packing applied to display ad allocation. In *ESA (1)*, volume 6346 of *Lecture Notes in Computer Science*, pages 182–194. Springer, 2010.
- [Kaneko and Nakamura, 1979] Mamoru Kaneko and Kenjiro Nakamura. The Nash social welfare function. *Econometrica: Journal of the Econometric Society*, pages 423–435, 1979.
- [Kesselheim *et al.*, 2014] Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. Primal beats dual on online packing lps in the random-order model. In *STOC*, pages 303–312. ACM, 2014.
- [Lee, 2017] Euiwoong Lee. APX-hardness of maximizing Nash social welfare with indivisible items. *Inf. Process. Lett.*, 122:17–20, 2017.
- [McGlaughlin and Garg, 2020] Peter McGlaughlin and Jugal Garg. Improving Nash social welfare approximations. *J. Artif. Intell. Res.*, 68:225–245, 2020.
- [Molinaro and Ravi, 2012] Marco Molinaro and R. Ravi. Geometry of online packing linear programs. In *ICALP (1)*, volume 7391 of *Lecture Notes in Computer Science*, pages 701–713. Springer, 2012.
- [Nash Jr, 1950] John F Nash Jr. The bargaining problem. *Econometrica: Journal of the econometric society*, pages 155–162, 1950.
- [Nguyen and Rothe, 2014] Trung Thanh Nguyen and Jörg Rothe. Minimizing envy and maximizing average Nash social welfare in the allocation of indivisible goods. *Discret. Appl. Math.*, 179:54–68, 2014.
- [Nguyen *et al.*, 2014] Nhan-Tam Nguyen, Trung Thanh Nguyen, Magnus Roos, and Jörg Rothe. Computational complexity and approximability of social welfare optimization in multiagent resource allocation. *Auton. Agents Multi Agent Syst.*, 28(2):256–289, 2014.
- [Ramezani and Endriss, 2009] Sara Ramezani and Ulle Endriss. Nash social welfare in multiagent resource allocation. In *AMEC/TADA*, volume 59 of *Lecture Notes in Business Information Processing*, pages 117–131. Springer, 2009.
- [Wu *et al.*, 2020] Xiaowei Wu, Bo Li, and Jiarui Gan. Budget-feasible maximum nash social welfare allocation is almost envy-free. *arXiv preprint arXiv:2012.03766*, 2020.