

The Long Arm of Nashian Allocation in Online p -Mean Welfare Maximization

Zhiyi Huang ✉ 

The University of Hong Kong, Pokfulam, Hong Kong

Chui Shan Lee ✉ 

The University of Hong Kong, Pokfulam, Hong Kong

Xinkai Shu¹ ✉ 

Max Planck Institute for Informatics, Saarbrücken, Germany

Zhaozi Wang² ✉ 

New York University, NY, US

Abstract

We study the online allocation of divisible items to n agents with additive valuations for p -mean welfare maximization, a problem introduced by Barman, Khan, and Maiti (2022). Our algorithmic and hardness results characterize the optimal competitive ratios for the entire spectrum of $-\infty \leq p \leq 1$. Surprisingly, our improved algorithms for all $p \leq 1/\log n$ are simply the greedy algorithm for the Nash welfare, supplemented with two auxiliary components to ensure all agents have non-zero utilities and to help a small number of agents with low utilities. In this sense, the long arm of Nashian allocation achieves near-optimal competitive ratios not only for Nash welfare but also all the way to egalitarian welfare.

2012 ACM Subject Classification Theory of computation → Online algorithms

Keywords and phrases Online Algorithms, Fair Division, Nash Welfare

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.98

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: <http://arxiv.org/abs/2504.13430>

Funding *Zhiyi Huang*: This work is supported by NSFC (No. 6212290003).

1 Introduction

Resource allocation is a fundamental challenge in many domains of Computer Science, Economics, and beyond, such as cloud computing (e.g., [21]) and food banks (e.g., [1]). The primary objectives in these scenarios are efficiency and fairness. For $-\infty \leq p \leq 1$, the p -means of the agents' utilities, known as the p -mean welfare, form an axiomatically justified family of objectives [31, Chapter 3] with different tradeoffs between these two factors. On one extreme when $p = 1$, this is the utilitarian welfare, i.e., the sum of the agents' utilities. On the other extreme when $p = -\infty$, this is the egalitarian welfare, i.e., the minimum utility among the agents. In the middle when $p = 0$ is the Nash welfare, which reconciles the two extremes and satisfies several notions of fairness including envy-freeness and proportionality [32].

This paper studies online resource allocation algorithms for maximizing the p -mean welfare among n agents when the items arrive sequentially and need to be allocated immediately upon arrival. This *online p -mean welfare maximization* problem was proposed by Barman et

¹ This work was done when the author was at the University of Hong Kong.

² This work was done when the author was at Shanghai Jiao Tong University.



© Zhiyi Huang, Chui Shan Lee, Xinkai Shu, and Zhaozi Wang;
licensed under Creative Commons License CC-BY 4.0

52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025).

Editors: Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis

Article No. 98; pp. 98:1–98:18



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



al. [8], assuming that 1) the items are divisible, 2) the agents' utilities are additive, and 3) each agent's utility for receiving all items, a.k.a., its monopolist utility, is equal to 1. For the whole spectrum of $-\infty \leq p \leq 1$, Barman et al. [8] gave competitive online algorithms and hardness results. However, there are gaps between their algorithmic and hardness results. For any negative constant p , the gap is polynomially large in the number of agents.

By contrast, the non-negative regime is better understood. For any positive p , the problem is equivalent to a special case of the *online matching with concave returns* problem introduced by Devanur and Jain [17]. Their online primal dual approach gave a $1/p$ competitive algorithm even without the third assumption above about unit monopolist utilities. We include an exposition of this algorithm and an improvement for $0 < p \leq 1/\log n$ using the third assumption in the full version of the paper. For Nash welfare ($p = 0$), Banerjee et al. [4] gave an $O(\log n)$ -competitive algorithm,³ which is the best possible.

1.1 Our Contributions

1.1.1 Optimal Competitive Ratios (up to Lower-Order Terms)

We improve the upper and lower bounds for the competitive ratios of online p -mean welfare maximization. These bounds characterize the optimal competitive ratios for all $-\infty \leq p \leq 1$. See Table 1 for a summary.

For $p \geq 1/\log n$, we prove that the $1/p$ competitive ratio obtained by Devanur and Jain [17] is optimal up to a $\log 1/p$ term even if the instance satisfies the assumption of unit monopolist utilities.

For the Nashian regime, i.e., when $-1/\log n \leq p \leq 1/\log n$, we improve the $O(\log^3 n)$ competitive ratio by Barman et al. [8] to $O(\log n)$, and give a corresponding hardness result matching it up to a doubly logarithmic factor. Note that for Nash welfare ($p = 0$), our result matches the optimal ratio of $O(\log n)$ obtained by Banerjee et al. [4].

Between the Nashian regime and the harmonic mean, i.e., when $-1 \leq p \leq -1/\log n$, we give an online algorithm that is $n^{|p|/(|p|+1)}$ -competitive, omitting lower-order terms. This is better than the state-of-the-art by Barman et al. [8] by a polynomial factor for $p \geq -1/4$, and a poly-logarithmic factor for other values of p . We complement the algorithmic result with a hard instance, showing that no online algorithm can do better up to a logarithmic term. This hardness result is a polynomial improvement compared to the existing $n^{|p|/(2|p|+1)}$ for any negative constant p .

Last but not least, from harmonic mean to egalitarian welfare, i.e., when $-\infty \leq p \leq -1$, we show that the optimal competitive ratio is \sqrt{n} up to a lower-order term. Our hardness result is a polynomial improvement on the existing $n^{|p|/(2|p|+1)}$ bound by Barman et al. [8] for any finite p ; on the other hand, our algorithmic result is slightly better than theirs by a $\sqrt{\log n}$ factor.

³ More precisely, their model assumes having predictions of the agents' monopolist utilities, which can be used to normalize the agents' utilities to restore the assumption of unit monopolist utilities. See also [24] for online algorithms that relaxed this assumption to allow the absence of predictions of agents' monopolist utilities but still require the max-min ratio of monopolist utilities to be bounded.

⁴ The omitted terms are of lower order for $p \leq -\omega(\log \log n / \log n)$. For $-O(\log \log n / \log n) \leq p \leq -1/\log n$, our results show that the optimal competitive ratio is poly-logarithmic, but do not characterize the degree of the poly-logarithms.

■ **Table 1** Summary of Results. We omit constant and lower-order terms for brevity, and defer the precise bounds to the subsequent sections. The bounds on some rows were known for a special value of p : Banerjee et al. [4] gave $\log n$ upper and lower bounds for $p = 0$, and Barman et al. [8] showed a \sqrt{n} lower bound for $p = -\infty$.

	Algorithmic Results		Hardness Results	
$1/\log n \leq p \leq 1$	$1/p$	Devanur and Jain [17]	$1/p$	Theorem 15
$-1/\log n \leq p \leq 1/\log n$	$\log n$	Corollary 4	$\log n$	Theorem 15
$-o(1) \leq p \leq -1/\log n^4$	$n^{ p }$	Theorem 3	$n^{ p }$	Theorem 16
$-1 \leq p \leq -\Omega(1)$	$n^{\frac{ p }{ p +1}}$	Theorem 6	$n^{\frac{ p }{ p +1}}$	Theorem 17
$-\infty \leq p \leq -1$	\sqrt{n}	Theorem 10; Barman et al. [8]	\sqrt{n}	Theorem 18

1.1.2 Long Arm of Nashian Allocation

Besides the improved competitive ratios, we find it conceptually interesting that our algorithmic results for the whole spectrum of $-\infty \leq p \leq 1/\log n$ are obtained by just two algorithms, both based on the greedy algorithm for the Nash welfare. *A priori*, it is surprising that the greedy algorithm for $p = 0$ also achieves nearly optimal competitive ratios for other values of p , even those that are far from zero. This can be viewed as further evidence for the effectiveness of Nash welfare maximization for balancing efficiency and fairness in the context of online optimization, echoing the unreasonable fairness from offline Nash welfare maximization showed by Caragiannis et al. [13].

More precisely, the improvements for $-1 \leq p \leq 1/\log n$ are obtained by combining the greedy algorithm for the Nash welfare with the common idea of distributing a constant fraction of each item uniformly to all agents to ensure an $\Omega(1/n)$ base utility for the agents (see, e.g., [4, 24], for previous analyses of this algorithm for $p = 0$). The base utilities allow us to avoid the irregularity of p -mean welfare when an agent has zero utility, for any $p \leq 0$.

We further address the cases of $-\infty \leq p \leq -1$ by introducing an auxiliary component that considers the other extreme of the spectrum, greedily maximizing the egalitarian welfare with respect to the agents' *regularized utilities*. An agent's regularized utility is the sum of its utility and a regularization term that is linear in the agent's monopolist utility for the remaining items. We stress that the greedy algorithm for the Nash welfare still plays the main role in this regime. In particular, it ensures that at most $\tilde{O}(\sqrt{n})$ agents may have low regularized utilities; the auxiliary component is designed to help these agents by making regularized egalitarian allocation.

1.1.3 Our Techniques

Next, we sketch the competitive analysis of the greedy algorithm for the Nash welfare. Consider the algorithm's allocation and any alternative allocation. Denote each agent a 's utilities for these allocations as U_a and \tilde{U}_a respectively. We will prove that the ratio of these utilities, i.e., \tilde{U}_a/U_a , is at most $O(\log n)$ averaging over all agents a . This is easy to prove in hindsight, based on 1) the greedy criteria with respect to the Nash welfare, and the fact that 2) the agents' maximum and minimum utilities differ by at most an $O(n)$ multiplicative factor due to the $\Omega(1/n)$ base utilities from distributing a constant fraction of each item uniformly to all agents. Nonetheless, this is powerful enough to derive all subsequent lemmas and competitive ratios in this paper; we call it the *Fundamental Lemma of Nashian Allocation* (Lemma 5).

In particular, we will use this fundamental lemma to derive upper bounds for the number of agents whose utilities are smaller than a threshold (Lemma 9), and the number of agents whose regularized utilities are smaller than a threshold (Lemma 12). Intuitively, these bounds are useful because the p -mean welfare for any negative $p < 0$ may be seen as a softmin function over the agents' utilities. Hence, lower bounding the p -mean welfare of the algorithm's allocation reduces to upper bounding the number of agents whose (regularized) utilities are too small.

Finally, our hardness results consider upper triangular instances that are prevalent in online resource allocation and online matching (e.g., [27, 26]). We supplement the upper triangular instances with items arriving at the end to satisfy the assumption of unit monopolist utilities.

1.2 Further Related Work

Online resource allocation problems have been studied extensively, including online packing [3, 12] and online matching [27, 30, 25]. Walsh [33] and Aleksandrov et al. [1] started the study of online fair division with agent and item arrivals respectively. See [2] for a survey.

Closest to this paper is the work by Barman et al. [8], which we have already discussed and compared against. Better results have been achieved in relaxed models of online allocation. Cohen and Panigrahi [16] studied the problem in a learning-augmented setting, i.e., where the algorithm has access to some extra (machine-learned) information. Hajiaghayi, Panigrahi, Khani, and Springer [23] considered the case of egalitarian welfare and indivisible items, assuming that the items arrive by a random order and the instance satisfies a large-market assumption.

Another line of work considered online resource allocation for other notions of fairness. Benade et al. [10] achieved envy-freeness in online allocation in the sense of no-regret learning. Zeng and Psomas [34] explored the fairness-efficiency tradeoffs in online fair division with respect to envy-freeness and Pareto optimality. Gkatzelis et al. [22] considered the online allocation of divisible items (to two agents) to satisfy envy-freeness and at the same time approximately maximize utilitarian welfare. Banerjee et al. [5] studied online allocation of public goods for proportional fairness, with predictions for the agents' monopolist utilities.

Finally, offline maximization for p -mean welfare has also been studied in a long line of work. Since the case of additive valuations and divisible items reduces to standard convex optimization, researchers focused on the harder case with indivisible items and non-additive valuations. For general $-\infty \leq p \leq 1$, Barman et al. [7] and Chaudhury et al. [15] independently gave $O(n)$ -approximation algorithms for maximizing the p -mean welfare when the valuations are subadditive.

For $p = 0$, maximizing the Nash welfare is APX-hard [28]. The state-of-the-art for additive valuations is an $e^{1/e}$ -approximation algorithm by Barman et al. [9], which has been generalized to the weighted case by Feng and Li [19]. Li and Vondrák [29] obtained a constant approximation for submodular valuations while Garg et al. [20] showed an $e^{1/(e-1)}$ lower bound. Recently, Dobzinski et al. [18] gave a constant approximation for subadditive valuations.

For $p = -\infty$, maximizing the egalitarian welfare for additive valuations and indivisible items, a.k.a., the Santa Claus problem (e.g., [6]), is also APX-hard [11]. Chakrabarty et al. [14] gave an $\tilde{O}(n^\varepsilon)$ -approximate algorithm that runs in $n^{O(1/\varepsilon)}$ time, for any $\varepsilon = \Omega(\log \log n / \log n)$.

2 Preliminaries

We write \log for the natural logarithm in this paper.

2.1 Model

2.1.1 Discrete Time

Consider the problem of online allocation of m *divisible* items I to n agents A . Agent a has value $v_{ai} \geq 0$ for item i . The items arrive one by one. The agents' values for an item are unknown initially and revealed when the item arrives. Upon an item's arrival, the algorithm must allocate it to some agent(s) immediately and irrevocably. Denote the algorithm's allocation as $x = (x_{ai})_{a \in A, i \in I}$, where $x_{ai} \geq 0$ is the portion of item i allocated to agent a . Agent a 's *utility* for allocation x is:

$$U_a = \sum_{i \in I} v_{ai} \cdot x_{ai} .$$

For some $p \leq 1$, we want to maximize the *p-mean welfare*:

$$\left(\frac{1}{n} \sum_{a \in A} U_a^p \right)^{\frac{1}{p}} . \quad (1)$$

When $p = 1$, this is the *utilitarian welfare* (up to a $1/n$ factor). When $p = 0$, Equation (1) is defined as the geometric mean of the agents' utilities, and is known as the *Nash welfare*. When $p = -1$, this becomes the harmonic mean of the agents' utilities; we therefore call it the *harmonic welfare*. Last but not least, when $p = -\infty$, Equation (1) is the minimum utility among the agents, and is called the *egalitarian welfare*.

2.1.2 Continuous Time

It will be more convenient to present our algorithms and analyses in the more general continuous-time model. Next, we present the model and then explain the reduction from the discrete-time model to the continuous-time model. Consider a continuum of infinitesimal items arriving in time horizon $[0, T)$ with unit arrival rate. We will refer to the item arriving at time t as item t , and thus, let $I = [0, T)$ denote the set of items. Agent a has unit value $v_a(t)$ for item t ; we assume that $v_a(t)$ is piecewise constant, changing its value only a finite number of times. Denote the algorithm's allocation as $x = (x_a(t))_{a \in A, t \in I}$ where $x_a(t)$ is the portion of item t allocated to agent a . Correspondingly, agent a 's utility for allocation x is:

$$U_a = \int_0^T v_a(t) x_a(t) dt .$$

Given any instance in the discrete-time model where we without loss of generality denote the set of m divisible items as $\{1, 2, \dots, m\}$, we can reinterpret it in the continuous-time model with $T = m$, such that agent a 's value for items $i - 1 \leq t < i$ is $v_a(t) = v_{ai}$, for any $1 \leq i \leq m$.

2.1.3 Assumption of Unit Monopolist Utilities

Following [8], we assume that the agents have unit monopolist utilities. That is, for any agent $a \in A$, we have:

$$\int_0^T v_a(t) dt = 1 .$$

Trivial hardness results exist if the algorithm has no information on the agents' monopolist utilities. Banerjee et al. [4] provided a hard instance that prevents any algorithm from performing better than $\Omega(n)$ -competitive in online Nash welfare maximization ($p = 0$). The proof of $\Omega(n)$ hardness can be generalized to any $p \leq 0$ using the same hard instance.

Our results still hold under a weaker assumption that the monopolist utilities are known and within a constant factor of each other (see the full version of the paper). This corresponds to having predictions for the agents' monopolist utilities, which could be derived from historical data and machine learning models. Moreover, the agents are of similar importance in the market. Further relaxing this assumption is an interesting research direction. See [4] and [24] for some related results for the Nash welfare.

2.2 Competitive Analysis

2.2.1 Offline Optimal Allocation

We will compare the p -mean welfare of the algorithm's allocation to the offline optimal benchmark, obtained from optimizing the allocation based on full knowledge of the instance. We can compute the offline optimal benchmark by solving a convex program:

$$\begin{aligned}
 & \text{maximize} && \left(\frac{1}{n} \sum_{a \in A} U_a^p \right)^{\frac{1}{p}} \\
 & \text{subject to} && \sum_{a \in A} x_a(t) \leq 1 && \forall t \in I \\
 & && U_a = \int_0^T v_a(t) x_a(t) dt && \forall a \in A \\
 & && x_a(t) \geq 0 && \forall a \in A, \forall t \in I
 \end{aligned}$$

We denote the optimal allocation as $x^* = (x_a^*(t))_{a \in A, t \in I}$ and the corresponding agents' utilities and p -mean welfare as $\{U_a^*\}_{a \in A}$ and OPT respectively. If there are multiple allocations that achieve the optimal p -mean welfare, we will pick an arbitrary one as x^* .

2.2.2 Competitive Online Algorithms

Denote the p -mean welfare for the online algorithm's allocation as ALG. An online algorithm is Γ -competitive if for every instance of online p -mean welfare maximization, we have:

$$\text{OPT} \leq \Gamma \cdot \text{ALG} .$$

2.3 Relaxation of Online p -Mean Welfare Maximization

We can guarantee a simple lower bound for all agents' utilities with the Uniform Allocation below.

Uniform Allocation

For each item $t \in I$, allocate it uniformly to all agents, i.e., $x_a(t) = 1/n$.

► **Lemma 1.** *Every agent $a \in A$ gets utility $1/n$ from Uniform Allocation.*

As a result of this simple bound, we can consider a relaxation of online p -mean welfare maximization, in which each agent starts with $1/n$ base utility for the algorithm's allocation. In other words, an agent a 's utility for an allocation x is:

$$U_a \stackrel{\text{def}}{=} \frac{1}{n} + \int_0^T v_a(t)x_a(t) dt .$$

We further write $U_a(t)$ for agent a 's utility for the items it received before time t , i.e.:

$$U_a(t) \stackrel{\text{def}}{=} \frac{1}{n} + \int_0^t v_a(s)x_a(s) ds .$$

As for the benchmark, we will still compare against the original optimal p -mean welfare without the $1/n$ base utility.

► **Lemma 2.** *If an online algorithm A is Γ -competitive for the relaxed online p -mean welfare maximization problem, then allocating half of each item by Uniform Allocation and the other half by algorithm A is 2Γ -competitive for the original problem.*

3 Nashian Regime: $-o(1) \leq p \leq 1/\log n$

This section considers the Nashian regime where p is close to zero. We will analyze the greedy algorithm for maximizing the Nash welfare and prove that it is competitive for all values of p in this regime simultaneously.

Recall that the Nash welfare is the case of $p = 0$, defined as the geometric mean of the agents' utilities. Equivalently, we may consider maximizing its logarithm (up to a factor n):

$$\sum_{a \in A} \log U_a .$$

For each item t , allocating it to agent a would increase this objective by:

$$\frac{v_a(t)}{U_a(t)} .$$

Hence, the Nashian Greedy algorithm allocates each item to an agent to maximize the above increment. We present below a formal definition of the algorithm.

Nashian Greedy

Initialization: Let $x_a(t) = 0$, $U_a(0) = \frac{1}{n}$ for any agent $a \in A$ and any time $t \in I$.

Online Decisions: Allocate each item $t \in I$ to an agent a that maximizes:

$$\frac{v_a(t)}{U_a(t)} ,$$

to greedily maximize the (logarithm of) Nash welfare:

$$\sum_{a \in A} \log U_a(t) .$$

The main results of this section are the following theorem and its corollary when $|p| \leq 1/\log n$.

► **Theorem 3.** For any $|p| = o(1)$, Nashian Greedy is $O(n^{|p|} \cdot \log n)$ -competitive for the relaxed online p -mean welfare maximization problem.

► **Corollary 4.** For any p such that $|p| \leq 1/\log n$, Nashian Greedy is $O(\log n)$ -competitive for the relaxed online p -mean welfare maximization problem.

Next, we present the most important lemma for the analysis of Nashian Greedy. Despite its simple form and proof, it is the foundation of all subsequent lemmas and competitive analyses in this paper.

► **Lemma 5 (Fundamental Lemma of Nashian Allocation).** Consider any allocation of the items $\tilde{x} = (\tilde{x}_a(t))_{a \in A, t \in I}$. For any time $t \in I$ and the agents' utilities for allocation \tilde{x} up to time t , denoted as $\tilde{U}(t) = (\tilde{U}_a(t))_{a \in A}$, we have:

$$\frac{1}{n} \sum_{a \in A} \frac{\tilde{U}_a(t)}{U_a(t)} \leq \log(n+1) .$$

Proof. Consider any item $s \in [0, t)$. Suppose that Nashian Greedy allocates it to agent a . By the definition of the algorithm, for any agent \tilde{a} we have:

$$\frac{v_a(s)}{U_a(s)} \geq \frac{v_{\tilde{a}}(s)}{U_{\tilde{a}}(s)} .$$

The left-hand-side equals the increment of the logarithm of Nash welfare for Nashian Greedy's allocation. Further, we relax the denominator of the right-hand-side to $U_{\tilde{a}}(t)$. We get that:

$$\frac{d}{ds} \sum_{a \in A} \log U_a(s) \geq \frac{v_{\tilde{a}}(s)}{U_{\tilde{a}}(t)} .$$

Multiplying this inequality by $\tilde{x}_a(s)$ and summing over all agents $\tilde{a} \in A$, we get that:

$$\frac{d}{ds} \sum_{a \in A} \log U_a(s) \geq \frac{d}{ds} \sum_{a \in A} \frac{\tilde{U}_a(s)}{U_a(t)}$$

Integrating over $s \in [0, t)$ gives:

$$\sum_{a \in A} \log \frac{U_a(t)}{U_a(0)} \geq \sum_{a \in A} \frac{\tilde{U}_a(t)}{U_a(t)} .$$

The lemma now follows by $1 + 1/n \geq U_a(t) \geq U_a(0) = 1/n$. ◀

Proof of Theorem 3. For $p = 0$, the theorem follows by considering Lemma 5 with $\tilde{x} = x^*$, i.e., the optimal allocation, and applying the AM-GM inequality to the left-hand-side.

Next, we consider the case when $p \neq 0$. Define an auxiliary allocation \tilde{x} that distributes half of each item uniformly to all agents, and the other half following the optimal allocation x^* . That is, for any agent $a \in A$ and any item $t \in I$:

$$\tilde{x}_a(t) \stackrel{\text{def}}{=} \frac{1}{2n} + \frac{1}{2}x_a^*(t) .$$

Let \tilde{U}_a be agent a 's utility for allocation \tilde{x} . By definition, we have:

$$\frac{1}{2n} \leq \tilde{U}_a \leq \frac{1}{2} + \frac{1}{2n} , \tag{2}$$

and also:

$$\left(\frac{1}{n} \sum_{a \in A} \tilde{U}_a^p\right)^{\frac{1}{p}} \geq \frac{\text{OPT}}{2}. \quad (3)$$

We write the p -mean welfare of Nashian Greedy's allocation as:

$$\text{ALG} = \left(\frac{1}{n} \sum_{a \in A} U_a^p\right)^{\frac{1}{p}} = \left(\frac{1}{n} \sum_{a \in A} \tilde{U}_a^p \cdot \left(\frac{\tilde{U}_a}{U_a}\right)^{-p}\right)^{\frac{1}{p}}. \quad (4)$$

Next, we introduce a set of auxiliary variables to denote:

$$z_a \stackrel{\text{def}}{=} \frac{\tilde{U}_a^p}{\sum_{a' \in A} \tilde{U}_{a'}^p}.$$

Comparing Equations (3) and (4), it suffices to show that:

$$\left(\sum_{a \in A} z_a \cdot \left(\frac{\tilde{U}_a}{U_a}\right)^{-p}\right)^{-\frac{1}{p}} \leq (n+1)^{|p|} \cdot \log(n+1).$$

By the definition of these auxiliary variables, we have $\sum_{a \in A} z_a = 1$. Hence, by applying the generalized mean inequality relating $(-p)$ -mean and 1-mean, where recall that $-p < 1$, we have:

$$\left(\sum_{a \in A} z_a \cdot \left(\frac{\tilde{U}_a}{U_a}\right)^{-p}\right)^{-\frac{1}{p}} \leq \sum_{a \in A} z_a \cdot \frac{\tilde{U}_a}{U_a}.$$

Further, by the range of \tilde{U}_a in Equation (2), we have:

$$z_a \leq \frac{1}{n} \left(\frac{\max_{a \in A} \tilde{U}_a}{\min_{a \in A} \tilde{U}_a}\right)^{|p|} \leq \frac{1}{n} \cdot (n+1)^{|p|}.$$

Combining the above with Lemma 5, we get that:

$$\left(\sum_{a \in A} z_a \cdot \left(\frac{\tilde{U}_a}{U_a}\right)^{-p}\right)^{-\frac{1}{p}} \leq (n+1)^{|p|} \cdot \frac{1}{n} \sum_{a \in A} \frac{\tilde{U}_a}{U_a} \leq (n+1)^{|p|} \cdot \log(n+1). \quad \blacktriangleleft$$

4 From Nash to Harmonic Welfare: $-1 \leq p \leq -\Omega(1)$

In this section, we continue to analyze the Nashian Greedy algorithm. Surprisingly, it remains competitive even when p is bounded away from zero. The resulting competitive ratios are nearly optimal for all values of p from Nash welfare to harmonic welfare.

► **Theorem 6.** *For any $-1 \leq p \leq -\Omega(1)$, Nashian Greedy is:*

$$O\left(n^{\frac{|p|}{|p|+1}} (\log n)^{\frac{1}{|p|+1}}\right)$$

competitive for the relaxed online p -mean welfare maximization problem.

4.1 Further Properties of Nashian Allocation

► **Definition 7** (Bad Agents). For any time $t \in I$ and any $\beta > 0$, let $B_\beta(t)$ be the set of β -bad agents whose utilities at time t are at most a β fraction of the optimal p -mean welfare, i.e.:

$$B_\beta(t) \stackrel{\text{def}}{=} \left\{ a \in A : U_a(t) \leq \beta \cdot \text{OPT} \right\} .$$

Further, we write B_β for $B_\beta(T)$, the set of β -bad agents at the end.

The next lemma considers a subset of β -bad agents, and upper bounds the sum of their utilities for the optimal allocation by a linear combination of $\beta \cdot \text{OPT}$, the upper bound of these agents' utilities for the algorithm's allocation, and the sum of these agents' remaining monopolist utilities.

► **Lemma 8.** For any time $t \in I$, any $\beta > 0$, and any subset of β -bad agents $S \subseteq B_\beta(t)$, we have:

$$\frac{1}{n} \sum_{a \in S} U_a^* \leq \beta \log(n+1) \cdot \text{OPT} + \frac{1}{n} \sum_{a \in S} \left(1 - \int_0^t v_a(s) \, ds \right) .$$

Proof. Recall that x^* denotes the optimal allocation. The left-hand-side of the lemma's inequality can be written as:

$$\frac{1}{n} \sum_{a \in S} \int_0^T v_a(s) x_a^*(s) \, ds .$$

On one hand, observe that:

$$\int_t^T v_a(s) x_a^*(s) \, ds \leq \int_t^T v_a(s) \, ds = 1 - \int_0^t v_a(s) \, ds .$$

On the other hand, by Lemma 5, we have:

$$\frac{1}{n} \sum_{a \in A} \frac{1}{U_a(t)} \int_0^t v_a(s) x_a^*(s) \, ds \leq \log(n+1) .$$

We now drop all agents outside subset S from the summation on the left-hand-side, and relax $U_a(t)$ to its upper bound $\beta \cdot \text{OPT}$ for the remaining β -bad agents $a \in S$. We get that:

$$\frac{1}{n} \sum_{a \in S} \int_0^t v_a(s) x_a^*(s) \, ds \leq \beta \log(n+1) \cdot \text{OPT} .$$

Putting together these two parts proves the lemma. ◀

► **Lemma 9.** For any $\beta > 0$, the fraction of β -bad agents at the end is at most:

$$\frac{|B_\beta|}{n} \leq (\beta \log(n+1))^{\frac{1}{p|+1}} .$$

Proof. By Lemma 8 with $t = T$ and $S = B_\beta$, we have:

$$\frac{1}{n} \sum_{a \in B_\beta} U_a^* \leq \beta \log(n+1) \cdot \text{OPT} .$$

Further, recall that $p < 0$. We have:

$$\text{OPT}^p = \frac{1}{n} \sum_{a \in A} (U_a^*)^{-|p|} \geq \frac{1}{n} \sum_{a \in B_\beta} (U_a^*)^{-|p|}.$$

Finally, by Hölder's inequality:

$$\left(\frac{1}{n} \sum_{a \in B_\beta} U_a^* \right)^{\frac{|p|}{|p|+1}} \left(\frac{1}{n} \sum_{a \in B_\beta} (U_a^*)^{-|p|} \right)^{\frac{1}{|p|+1}} \geq \frac{1}{n} \sum_{a \in B_\beta} 1 = \frac{|B_\beta|}{n}.$$

Combining these inequalities proves the lemma. \blacktriangleleft

4.2 Proof of Theorem 6

We compare the p -mean welfare of Nashian Greedy to the optimal benchmark by:

$$\left(\frac{\text{ALG}}{\text{OPT}} \right)^p = \frac{1}{n} \sum_{a \in A} \left(\frac{U_a}{\text{OPT}} \right)^p.$$

By $U_a \geq 1/n$ and $\text{OPT} \leq 1$, and recalling that $p < 0$, we have :

$$\left(\frac{U_a}{\text{OPT}} \right)^p \leq n^{|p|}.$$

Therefore, the above ratio can be written as:

$$\begin{aligned} \left(\frac{\text{ALG}}{\text{OPT}} \right)^p &= \int_0^{n^{|p|}} (\text{fraction of agents with } (U_a/\text{OPT})^p \geq \alpha) \, d\alpha \\ &\leq \int_0^{n^{|p|}} (\alpha^{\frac{1}{p}} \log(n+1))^{\frac{|p|}{|p|+1}} \, d\alpha && \text{(Lemma 9)} \\ &= \int_0^{n^{|p|}} (\log(n+1))^{\frac{|p|}{|p|+1}} \alpha^{-\frac{1}{|p|+1}} \, d\alpha \\ &= \frac{|p|+1}{|p|} \cdot n^{\frac{|p|^2}{|p|+1}} (\log(n+1))^{\frac{|p|}{|p|+1}}. \end{aligned}$$

Taking p -th root on both sides gives:

$$\frac{\text{ALG}}{\text{OPT}} \geq \underbrace{\left(\frac{|p|}{|p|+1} \right)^{\frac{1}{|p|}}}_{\Omega(1) \text{ for } |p| = \Omega(1)} \cdot n^{-\frac{|p|}{|p|+1}} (\log(n+1))^{-\frac{1}{|p|+1}}.$$

5 From Harmonic to Egalitarian Welfare: $-\infty \leq p \leq -1$

This section considers the regime between harmonic welfare and egalitarian welfare. In this case, we need to introduce an auxiliary component to obtain a new lower bound of the agents' utilities that is better than the $1/n$ base utility given by Uniform Allocation.

Although Nashian Greedy on its own fails to provide a better bound, we observe that the number of agents who are in trouble cannot exceed $\sqrt{n \log(n+1)}$ (see Corollary 13). Hence, we may reserve a constant fraction of each item to help these agents.

How can we identify the agents in trouble? Naïvely, it is natural to consider an egalitarian greedy algorithm that allocates the item to the agents who currently have the lowest utilities. However, an agent may have low utility just because its valuable items are yet to arrive. It

would be a mistake to allocate items to such an agent purely based on egalitarian consideration. Instead, we introduce for each agent $a \in A$ a regularization term R_a , which equals the agent's remaining monopolist utility, i.e., its total utility for the items yet to arrived, scaled by a $\sqrt{n \log(n+1)}$ factor that is driven by the analysis. Our new component is the egalitarian greedy allocation with respect to the regularized utilities $U_a + R_a$.

We remark that Barman et al. [8] also introduced an algorithmic sub-routine to identify the agents in trouble, and already observed the necessity to consider both the agents' utilities and their remaining monopolist utilities. Our regularized egalitarian allocation gives an alternative, and in our opinion, more natural way to implement this idea. The resulting competitive ratios are also asymptotically better than those in [8].

Below we present the formal definition of this Mixed Greedy algorithm that combines Nashian Greedy and Regularized Egalitarian Greedy. For cleaner constants in the analysis, we further relax the problem by letting there be two copies of each item. This relaxation only affects the competitive ratio by a constant factor.

Mixed Greedy (Nashian Greedy and Regularized Egalitarian Greedy)

Initialization: Let $x_a(t) = 0$, $U_a(0) = \frac{1}{n}$ for any agent $a \in A$ and any time $t \in I$.

Regularization: Define the regularizer of any agent $a \in A$ as:

$$R_a(t) = \frac{1}{\Phi} \left(1 - \int_0^t v_a(s) ds \right),$$

where $\Phi = \sqrt{n \log(n+1)}$.

Online Decisions: For each item $t \in I$, which comes in two copies:

1. Allocate the *Nashian copy* using Nashian Greedy.
2. Allocate the *egalitarian copy* to the agents with the smallest regularized utility:

$$U_a(t) + R_a(t),$$

to greedily maximize the regularized egalitarian welfare:

$$\min_{a \in A} (U_a(t) + R_a(t)).$$

This algorithm achieves the same and nearly optimal competitive ratio simultaneously for all values of p from harmonic welfare to egalitarian welfare.

► **Theorem 10.** *For any $-\infty \leq p \leq -1$, Mixed Greedy is $O(\sqrt{n \log n})$ -competitive for the relaxed online p -mean welfare maximization problem.*

5.1 Properties of the Regularized Utilities

► **Definition 11 (Critical Agents).** *For any time $t \in I$, and for any $\beta \geq 0$, let $C_\beta(t)$ be the set of β -critical agents at time t , whose regularized utilities are at most a β fraction of the optimal p -mean welfare, i.e.:*

$$C_\beta(t) = \left\{ a \in A : U_a(t) + R_a(t) \leq \beta \cdot \text{OPT} \right\}.$$

► **Lemma 12.** *For any time $t \in I$ and any $\beta > 0$, the fraction of β -critical agents at time t is upper bounded by:*

$$\frac{|C_\beta(t)|}{n} \leq \max \left\{ (2\beta \log(n+1))^{\frac{|p|}{|p|+1}}, (2\Phi\beta)^{|p|} \right\}.$$

Proof. On one hand, by the definition of p -mean welfare and $p < 0$, we have:

$$\text{OPT}^p = \frac{1}{n} \sum_{a \in A} (U_a^*)^{-|p|} \geq \frac{1}{n} \sum_{a \in C_\beta(t)} (U_a^*)^{-|p|}.$$

On the other hand, all β -critical agents are β -bad. We have:

$$\begin{aligned} \frac{1}{n} \sum_{a \in C_\beta(t)} U_a^* &\leq \beta \log(n+1) \cdot \text{OPT} + \frac{1}{n} \sum_{a \in C_\beta(t)} \left(1 - \int_0^t v_a(s) ds\right) && \text{(Lemma 8)} \\ &= \beta \log(n+1) \cdot \text{OPT} + \frac{1}{n} \sum_{a \in C_\beta(t)} \Phi \cdot R_a(t) \\ &\leq \left(\beta \log(n+1) + \frac{|C_\beta(t)|}{n} \cdot \Phi\beta \right) \cdot \text{OPT}. \end{aligned}$$

$(R_a(t) \leq U_a(t) + R_a(t) \leq \beta \cdot \text{OPT})$

By Hölder's inequality:

$$\left(\frac{1}{n} \sum_{a \in C_\beta(t)} (U_a^*)^{-|p|} \right)^{\frac{1}{|p|+1}} \left(\frac{1}{n} \sum_{a \in C_\beta(t)} U_a^* \right)^{\frac{|p|}{|p|+1}} \geq \frac{1}{n} \sum_{a \in C_\beta(t)} 1 = \frac{|C_\beta(t)|}{n}.$$

Combining these inequalities gives:

$$\left(\beta \log(n+1) + \frac{|C_\beta(t)|}{n} \cdot \Phi\beta \right)^{\frac{|p|}{|p|+1}} \geq \frac{|C_\beta(t)|}{n}.$$

Rearranging terms, we have:

$$\beta \log(n+1) \cdot \left(\frac{n}{|C_\beta(t)|} \right)^{\frac{|p|+1}{|p|}} + \Phi\beta \cdot \left(\frac{n}{|C_\beta(t)|} \right)^{\frac{1}{|p|}} \geq 1.$$

Hence, either the first part is at least a half, in which case:

$$\frac{|C_\beta(t)|}{n} \leq (2\beta \log(n+1))^{\frac{|p|}{|p|+1}},$$

or the second part is at least a half, in which case:

$$\frac{|C_\beta(t)|}{n} \leq (2\Phi\beta)^{|p|}.$$

In either case, the lemma follows. ◀

► **Corollary 13.** *For any time $t \in I$, and:*

$$\beta^* = \frac{1}{2} \cdot n^{-\frac{1}{2} - \frac{1}{2|p|}} (\log(n+1))^{-\frac{1}{2} + \frac{1}{2|p|}}, \quad (5)$$

we have:

$$|C_{\beta^*}(t)| \leq \sqrt{n \log(n+1)}.$$

98:14 The Long Arm of Nashian Allocation in Online p -Mean Welfare Maximization

Using Corollary 13, we derive a universal lower bound for all agents' utilities.

► **Lemma 14.** *For the choice of β^* in Equation (5), and any agent $a \in A$, we have:*

$$U_a \geq \beta^* \cdot \text{OPT} .$$

Proof. We will prove a stronger claim that for any agent $a \in A$ and any time $t \in I$:

$$U_a(t) + R_a(t) \geq \beta^* \cdot \text{OPT} . \tag{6}$$

Then, the lemma holds as the special case when $t = T$ because $R_a(T) = 0$.

Initially at time $t = 0$, we have:

$$R_a(0) = \frac{1}{\Phi} > \beta^* \geq \beta^* \cdot \text{OPT} .$$

To prove that Equation (6) holds at all time t , it suffices to show that for any time $t \in I$ when there is at least one β^* -critical agent $a \in C_{\beta^*}(t)$, the allocation of the egalitarian copy of item t weakly increases the regularized egalitarian welfare.

For any critical agent $a \in C_{\beta^*}(t)$, we have:

$$\frac{d}{dt} R_a(t) = -\frac{v_a(t)}{\Phi} .$$

Further, Corollary 13 asserts that at most $\sqrt{n \log(n+1)}$ agents are β^* -critical. Hence, allocating the egalitarian copy of item t equally among these β^* -critical agents would have yielded:

$$\frac{d}{dt} U_a(t) \geq \frac{v_a(t)}{\sqrt{n \log(n+1)}} = \frac{v_a(t)}{\Phi} ,$$

and weakly increased the regularized egalitarian welfare. The greedy allocation of the algorithm would only do better. ◀

5.2 Proof of Theorem 10

The proof is almost verbatim to that of Theorem 6, except that we will use the newly developed Lemma 14 to lower bound the agents' utilities, replacing the basic bound $U_a \geq 1/n$.

Consider the p -th power of the Mixed Greedy algorithm's p -mean welfare, normalized by the p -th power of OPT:

$$\left(\frac{\text{ALG}}{\text{OPT}} \right)^p = \frac{1}{n} \sum_{a \in A} \left(\frac{U_a}{\text{OPT}} \right)^p .$$

By Lemma 14 and $p < 0$, we have:

$$\left(\frac{U_a}{\text{OPT}} \right)^p \leq (\beta^*)^p .$$

The above ratio can therefore be written as:

$$\begin{aligned}
\left(\frac{\text{ALG}}{\text{OPT}}\right)^p &= \int_0^{(\beta^*)^p} (\text{fraction of agents with } (U_a/\text{OPT})^p \geq \alpha) \, d\alpha \\
&\leq \int_0^{(\beta^*)^p} (\alpha^{\frac{1}{p}} \log(n+1))^{\frac{|p|}{|p|+1}} \, d\alpha && \text{(Lemma 9)} \\
&= \int_0^{(\beta^*)^p} (\log(n+1))^{\frac{|p|}{|p|+1}} \alpha^{-\frac{1}{|p|+1}} \, d\alpha \\
&= \frac{|p|+1}{|p|} (\beta^*)^{-\frac{|p|^2}{|p|+1}} (\log(n+1))^{\frac{|p|}{|p|+1}} \\
&= \frac{|p|+1}{|p|} 2^{\frac{|p|^2}{|p|+1}} (n \log(n+1))^{\frac{|p|}{2}}.
\end{aligned}$$

Taking p -th root on both sides gives:

$$\frac{\text{ALG}}{\text{OPT}} \geq \underbrace{\left(\frac{|p|}{|p|+1}\right)^{\frac{1}{|p|}} 2^{-\frac{|p|}{|p|+1}}}_{\Omega(1) \text{ for } p \leq -1} \cdot \left(\sqrt{n \log(n+1)}\right)^{-1}.$$

6 Hardness Results

In this section, we complement our algorithmic results with nearly-tight hardness results. The proofs are deferred to the full version of the paper.

6.1 Hardness for the Nashian and Positive Regimes: $p \geq -1/\log n$

Recall that the online primal dual algorithm by Devanur and Jain [17] is $1/p$ -competitive for $0 < p \leq 1$, and Nashian Greedy is $O(\log n)$ -competitive for $-1/\log n \leq p \leq 1/\log n$ (Corollary 4). Theorem 15 complement these competitive ratios with hardness results that match them up to a lower-order term.

► **Theorem 15.** *For any $p \leq 1$ and any online algorithm for online p -mean welfare maximization, the competitive ratio is no smaller than:*

$$\begin{cases} \Omega\left(\frac{1}{p}\right) & \text{if } p \geq \frac{\log \log n}{\log n}; \\ \Omega\left(\frac{\log n}{\log \log n}\right) & \text{otherwise,} \end{cases}$$

even when the agents have binary valuations (and unit monopolist utilities).

This asymptotic lower bound holds for a sufficiently large number of agents. It implicitly covers three cases. If p is a positive constant independent of the number of agents n , it characterizes how the competitive ratio changes as p tends to 0. If $p(n)$ is a function of the number of agents n that converges to zero as n goes to infinity, e.g., $p(n) = 1/\log n$, the theorem asserts that for the competitive ratio is no smaller than $\Omega(1/p(n))$ if $p(n)$ converges to zero slower than function $\log \log n / \log n$, or no smaller than $\Omega(\log n / \log \log n)$ otherwise. The latter bound also applies when p is a negative constant or a negative function $p(n)$ that converges to zero.

6.2 Hardness for the Negative Regime: $p \leq -1/\log n$

Recall that Nashian Greedy is $n^{|p|}$ -competitive for $-o(1) \leq p \leq -1/\log n$ and $n^{|p|/(|p|+1)}$ -competitive for $-1 \leq p \leq -\Omega(1)$, up to lower-order and logarithmic factors (Theorem 3). Moreover, Mixed Greedy is $O(\sqrt{n \log n})$ -competitive for $-\infty \leq p \leq -1$ (Theorem 10). Theorems 16–18 complement these algorithmic results with almost matching lower bounds.

► **Theorem 16.** *For any $-o(1) \leq p \leq -\omega(\log \log n / \log n)$ and any online algorithm for online p -mean welfare maximization, the competitive ratio is no smaller than $n^{|p|-o(|p|)}$, even when the agents have binary valuations (and unit monopolist utilities).*

► **Theorem 17.** *For any $-1 \leq p \leq \Omega(1)$ and any online algorithm for online p -mean welfare maximization, the competitive ratio is no smaller than $n^{\frac{|p|}{|p|+1}-o(1)}$, even when the agents have binary valuations (and unit monopolist utilities).*

► **Theorem 18.** *For any $-\infty \leq p \leq -1$ and any online algorithm for online p -mean welfare maximization, the competitive ratio is no smaller than $n^{\frac{1}{2}-o(1)}$, even when the agents have binary valuations (and unit monopolist utilities).*

References

- 1 Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh. Online fair division: analysing a food bank problem. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, pages 2540–2546, 2015. doi:10.5555/2832581.2832604.
- 2 Martin Aleksandrov and Toby Walsh. Online fair division: A survey. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, pages 13557–13562, 2020. doi:10.1609/aaai.v34i09.7081.
- 3 Noga Alon, Baruch Awerbuch, Yossi Azar, Niv Buchbinder, and Joseph Naor. A general approach to online network optimization problems. *ACM Transactions on Algorithms*, 2(4):640–660, 2006. doi:10.1145/1198513.1198522.
- 4 Siddhartha Banerjee, Vasilis Gkatzelis, Artur Gorokh, and Billy Jin. Online Nash social welfare maximization with predictions. In *Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1–19. SIAM, 2022. doi:10.1137/1.9781611977073.1.
- 5 Siddhartha Banerjee, Vasilis Gkatzelis, Safwan Hossain, Billy Jin, Evi Micha, and Nisarg Shah. Proportionally fair online allocation of public goods with predictions. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence*, pages 20–28, 2023. doi:10.24963/ijcai.2023/3.
- 6 Nikhil Bansal and Maxim Sviridenko. The santa claus problem. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 31–40, 2006. doi:10.1145/1132516.1132522.
- 7 Siddharth Barman, Umang Bhaskar, Anand Krishna, and Ranjani G Sundaram. Tight approximation algorithms for p -mean welfare under subadditive valuations. In *Proceedings of the 28th Annual European Symposium on Algorithms*, pages 11:1–11:17, 2020. doi:10.4230/LIPIcs.ESA.2020.11.
- 8 Siddharth Barman, Arindam Khan, and Arnab Maiti. Universal and tight online algorithms for generalized-mean welfare. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence*, pages 4793–4800, 2022. doi:10.1609/aaai.v36i5.20406.
- 9 Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation*, pages 557–574, 2018. doi:10.1145/3219166.3219176.
- 10 Gerdus Benadè, Aleksandr M Kazachkov, Ariel D Procaccia, and Christos-Alexandros Psomas. How to make envy vanish over time. In *Proceedings of the 19th ACM Conference on Economics and Computation*, pages 593–610, 2018. doi:10.1145/3219166.3219179.

- 11 Ivona Bezáková and Varsha Dani. Allocating indivisible goods. *ACM SIGecom Exchanges*, 5(3):11–18, 2005. doi:10.1145/1120680.1120683.
- 12 Niv Buchbinder and Joseph Naor. Online primal-dual algorithms for covering and packing. *Mathematics of Operations Research*, 34(2):270–286, 2009. doi:10.1287/moor.1080.0363.
- 13 Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. *ACM Transactions on Economics and Computation*, 7(3):1–32, 2019. doi:10.1145/3355902.
- 14 Deeparnab Chakrabarty, Julia Chuzhoy, and Sanjeev Khanna. On allocating goods to maximize fairness. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 107–116, 2009. doi:10.1109/FOCS.2009.51.
- 15 Bhaskar Ray Chaudhury, Jugal Garg, and Ruta Mehta. Fair and efficient allocations under subadditive valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, pages 5269–5276, 2021. doi:10.1609/aaai.v35i6.16665.
- 16 Ilan Reuven Cohen and Debmalya Panigrahi. A general framework for learning-augmented online allocation. In *Proceedings of the 50th International Colloquium on Automata, Languages, and Programming*, pages 43:1–43:21, 2023. doi:10.4230/LIPIcs.ICALP.2023.43.
- 17 Nikhil R Devanur and Kamal Jain. Online matching with concave returns. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, pages 137–144, 2012. doi:10.1145/2213977.2213992.
- 18 Shahar Dobzinski, Wenzheng Li, Aviad Rubinfeld, and Jan Vondrák. A constant-factor approximation for nash social welfare with subadditive valuations. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 467–478, 2024. doi:10.1145/3618260.3649740.
- 19 Yuda Feng and Shi Li. A note on approximating weighted nash social welfare with additive valuations. In *Proceedings of the 51st EATCS International Colloquium on Automata, Languages and Programming*, 2024. doi:10.4230/LIPIcs.ICALP.2024.63.
- 20 Jugal Garg, Pooja Kulkarni, and Rucha Kulkarni. Approximating Nash social welfare under submodular valuations through (un)matchings. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2673–2687, 2020. doi:10.1145/3613452.
- 21 Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Symposium on Networked Systems Design and Implementation*, 2011. doi:10.5555/1972457.1972490.
- 22 Vasilis Gkatzelis, Alexandros Psomas, and Xizhi Tan. Fair and efficient online allocations with normalized valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, pages 5440–5447, 2021. doi:10.1609/aaai.v35i6.16685.
- 23 MohammadTaghi Hajiaghayi, Debmalya Panigrahi, Mohammad Khani, and Max Springer. Online algorithms for the Santa Claus problem. In *Proceedings of the 36th Annual Conference on Advances in Neural Information Processing Systems*, pages 30732–30743, 2022. doi:10.5555/3600270.3602498.
- 24 Zhiyi Huang, Minming Li, Xinkai Shu, and Tianze Wei. Online nash welfare maximization without predictions. In *Proceedings of the 19th International Conference on Web and Internet Economics*, pages 402–419. Springer, 2023. doi:10.1007/978-3-031-48974-7_23.
- 25 Zhiyi Huang, Zhihao Gavin Tang, and David Wajc. Online matching: A brief survey. *ACM SIGecom Exchanges*, 22:135–158, 2024. doi:10.1145/3699824.3699837.
- 26 Bala Kalyanasundaram and Kirk R Pruhs. An optimal deterministic algorithm for on-line b-matching. *Theoretical Computer Science*, 233(1-2):319–325, 2000. doi:10.1016/S0304-3975(99)00140-1.
- 27 Richard Karp, Umesh Vazirani, and Vijay Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*, pages 352–358, 1990. doi:10.1145/100216.100262.

- 28 Euiwoong Lee. APX-hardness of maximizing Nash social welfare with indivisible items. *Information Processing Letters*, 122:17–20, 2017. doi:10.1016/j.ip1.2017.01.012.
- 29 Wenzheng Li and Jan Vondrák. A constant-factor approximation algorithm for nash social welfare with submodular valuations. In *Proceedings of the 62nd Annual IEEE Symposium on Foundations of Computer Science*, pages 25–36, 2022. doi:10.1109/FOCS52979.2021.00012.
- 30 Aranyak Mehta. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science*, 8(4):265–368, 2013. doi:10.1561/04000000057.
- 31 Hervé Moulin. *Fair division and collective welfare*. MIT press, 2004. doi:10.7551/mitpress/2954.001.0001.
- 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 Toby Walsh. Online cake cutting. In *Proceedings of the 2nd International Conference on Algorithmic Decision Theory*, pages 292–305. Springer, 2011. doi:10.1007/978-3-642-24873-3_22.
- 34 David Zeng and Alexandros Psomas. Fairness-efficiency tradeoffs in dynamic fair division. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 911–912, 2020. doi:10.1145/3391403.3399467.