

Online Nash Welfare Maximization Without Predictions ^{*}

Zhiyi Huang [†] Minming Li [‡] Xinkai Shu [†] Tianze Wei [‡]

February 2023

Abstract

The maximization of Nash welfare, which equals the geometric mean of agents' utilities, is widely studied because it balances efficiency and fairness in resource allocation problems. Banerjee, Gkatzelis, Gorokh, and Jin (2022) recently introduced the model of online Nash welfare maximization for T divisible items and N agents with additive utilities with predictions of each agent's utility for receiving all items. They gave online algorithms whose competitive ratios are logarithmic. We initiate the study of online Nash welfare maximization *without predictions*, assuming either that the agents' utilities for receiving all items differ by a bounded ratio, or that their utilities for the Nash welfare maximizing allocation differ by a bounded ratio. We design online algorithms whose competitive ratios only depend on the logarithms of the aforementioned ratios of agents' utilities and the number of agents.

^{*}This is the second version of the paper on arXiv, which improves exposition based on feedbacks that we received after posting the first version. The main results stay the same. We thank anonymous reviewers for insightful and constructive comments, and thank Arindam Khan for pointing out several inaccurate references to previous work.

[†]University of Hong Kong. Email: zhiyi@cs.hku.hk, xkshu@cs.hku.hk.

[‡]City University of Hong Kong. Email: minming.li@cityu.edu.hk, t.z.wei-8@my.cityu.edu.hk.

1 Introduction

Suppose that you run a daycare center. From time to time, some organizations donate candies for children who have different preferences over various kinds of candies: Alice likes chocolates, Bob prefers gummy bears, Charlie favors lollipops, and so forth. The candies come in abundance, so you may consider it as a resource allocation problem of *divisible* items. How would you distribute the candies to the children?

Naturally, we would like to allocate the candies based on the children’s values for them. We will assume that the children have *additive utilities* for receiving bundles of candies. Following this idea, it may be tempting to allocate the candies to the children in a way that maximizes the sum of children’s utilities for the candies that they received, a.k.a., the social welfare. The example below, however, highlights the potential unfairness of this approach.

Example 1. Consider distributing two packs of chocolates and two packs of gummy bears to two children Alice and Bob. Alice has values 100 and 15 for receiving a pack of chocolates and gummy bears respectively. Bob, on the other hand, has values 1 and 10 for chocolates and gummy bears respectively.

The social welfare maximizing allocation gives everything to Alice for her higher values, and leaves nothing for Bob. This is blatantly unfair, especially at a daycare center.

What if we distribute evenly and let each child have a pack of chocolates and a pack of gummy bears? While this is fair, it is inefficient. The gummy bears contribute little to Alice’s utility, while the chocolates do not add much to Bob’s utility.

The question then becomes how to allocate the candies if we want to strike a balance between efficiency and fairness. In the above example, allocating both packs of chocolates to Alice and both packs of gummy bears to Bob is a decent option. In general, maximizing the *Nash welfare*, i.e., the geometric mean of the children’s utilities, is a good proxy for balancing efficiency and fairness. Indeed, the allocation that maximizes Nash welfare also satisfies several standard notions of fairness, such as envy-freeness [11] and proportionality [30]. It also coincides with the aforementioned decent allocation in the above example. Further, the Nash welfare maximization problem is *scale invariant*: if we increase an agent’s values for all items by the same multiplicative factor, the Nash welfare maximizing allocation would remain the same. Moreover, maximizing the Nash welfare with divisible items and additive utilities reduces to solving the Eisenberg-Gale convex program [16]. We can therefore compute this allocation in polynomial time.

There is one more complication though. We cannot predict when the donations of candies will be made, what kinds of candies will be donated, and in what quantities. When some organizations make a donation, we need to distribute the candies to the children immediately without much information about future donations.

Besides the toy example above, many real online resource allocation problems for heterogeneous agents require balance between fairness and efficiency, e.g., allocating food each day to people in need [29], or allocating shared computational resources to users [20]. Hence, our task becomes designing online allocation algorithms for Nash welfare maximization, whose performance is evaluated by *competitive ratio*, that is, the worst ratio between the online algorithm’s solution and the offline optimal solution. From now on, we will more generally talk about items and agents.

This problem was first studied by Banerjee, Gkatzelis, Gorokh, and Jin [5]. They noticed that under worst-case analysis, no algorithm can have a competitive ratio better than the trivial $O(N)$, where N is the number of agents. They then turned to the model of online algorithms with predictions, in which the online algorithm is further given some predictions of each agent’s utility for receiving the set of all items, which they called the *monopolist utility*. If the predictions

Table 1: Summary of results. Here λ^* and μ^* are the smallest numbers for which the instance is λ^* -balanced and μ^* -impartial. We omit the constants and $\log \log$ factors for brevity. We also omit the trivial bound of N .

	Algorithm	Lower Bound
λ^* -Balanced	$\log \lambda^* N$	$\log \lambda^* N$
μ^* -Impartial	$\min \{ \log \mu^* N, \log^2 \mu^* \}$	$\log \mu^*$
(α, β) -Predictions	$\alpha \log \beta N$ [5] \rightarrow $\log \alpha \beta N$	

were accurate, their algorithm would be $O(\log N)$ -competitive.¹ They further proved that the competitive ratio is asymptotically optimal for the problem.

1.1 Our Contribution

Coming up with accurate predictions is not an easy task at all. Is it truly hopeless to design good online algorithms *without predictions*? The hard instance which Banerjee et al. [5] designed to rule out the existence of good online algorithms without predictions requires the agents' values to differ by an exponential factor, a rare phenomenon in real resource allocation problems. Can we design online algorithms for "natural instances" where agents' values are not as extreme?

We introduce the notions of *balanced* and *impartial* instances. An instance is λ -balanced if the agents' monopolist utilities differ by at most a multiplicative factor λ . This is closely related to the model of online algorithms with predictions by Banerjee et al. [5]. If we had accurate predictions of the agents' monopolist utilities, we could then normalize the agents' values accordingly to obtain a 1-balanced instance using the fact that Nash welfare maximization is scale invariant. If we had (α, β) -approximate predictions, i.e., if they were at worst an α -factor larger or a β -factor smaller than the agents' monopolist utilities, the normalized instance would be $\alpha\beta$ -balanced. Therefore, any competitive online algorithm for balanced instances further implies a competitive online algorithm with predictions.

Our main result for λ -balanced instances is a $\log^{1+o(1)}(\lambda N)$ -competitive online algorithm. We further show that this competitive ratio is nearly optimal in the sense that no online algorithm can be better than $\log^{1-o(1)}(\lambda N)$ competitive. By the aforementioned reduction, this also implies a $\log^{1+o(1)}(\alpha\beta N)$ -competitive online algorithm for the model with (α, β) -approximate predictions, which improves the $O(\alpha \log \beta N)$ -competitive online algorithm given by Banerjee et al. [5] when α is sufficiently large.

On the other hand, we say that an instance is μ -impartial if the agents' utilities for the Nash welfare maximizing allocation differ by at most a multiplicative factor μ . For μ -impartial instances, our main result is a $\log^{2+o(1)} \mu$ -competitive online algorithm. We remark that this competitive ratio is independent of the number of agents or items, and thus this would be an $O(1)$ -competitive algorithm if the instance is $O(1)$ -impartial. In the special case when the agents' values for different items are binary, we have an improved competitive ratio $\log^{1+o(1)} \mu$; we further prove that it is nearly optimal.

See Table 1 for a summary of our results.

¹Banerjee et al. [5] also gave an $O(\log T)$ -competitive ratio where T is the number of items, under the assumption that items have unit supplies. Since this paper considers arbitrary supplies, the dependence in T is no longer valid.

1.2 Other Related Works

There is a vast literature on online algorithms for resource allocation problems, such as online packing (c.f., Alon et al. [1] and Buchbinder and Naor [10]) and online matching problems (c.f., Karp, Vazirani, and Vazirani [23] and Mehta [26]). Most related to this paper is the work by Devanur and Jain [15] on online matching with concave returns. Maximizing the Nash welfare can be equivalently formulated as maximizing the sum of logarithms of the agents' utilities, a special case of the model of Devanur and Jain [15] when the return functions are the logarithm. However, their results do not imply competitive online algorithms for our problem because they studied the multiplicative competitive ratio with respect to (w.r.t.) the sum of log-utilities. By contrast, Banerjee et al. [5] and this paper consider the multiplicative competitive ratio w.r.t. the geometric mean of utilities, which corresponds to the additive approximation factor w.r.t. the sum of log-utilities. Barman et al. [9] studied online algorithms for the more general p -mean welfare maximization which captures the Nash welfare as a special case when $p = 0$. Their result for Nash welfare can be interpreted as an $O(\log^3 N)$ -competitive algorithm for 1-balanced instances; by contrast our algorithm is $O(\log^{1+o(1)}(\lambda N))$ -competitive for general λ -balanced instances.

Nash welfare was introduced about 70 years ago [22, 28]. Its maximization has been extensively studied in the offline setting. The problem of divisible items and additive utilities is equivalent to the Eisenberg-Gale convex program [16]. Much effort has been devoted to the problem of indivisible items and various kinds of utilities. Even for additive utilities, the problem is NP-hard [27] and APX-hard [24]. Cole and Gkatzelis [13] gave the first constant-approximation algorithm for additive utilities. Subsequently, Anari et al. [2] applied the theory of stable polynomials to get an ϵ -approximation algorithm. Cole et al. [14] further improves the analysis to show a 2-approximation. Finally, Barman, Krishnamurthy, and Vaish [7] gave a new algorithm that achieves the state-of-the-art $e^{1/e} < 1.45$ -approximation. The problem with more general utilities has also been studied. Garg, Hoefer, and Mehlhorn [17] designed a constant-approximation algorithm for budget-additive utilities. Anari et al. [3] introduced a constant-approximation algorithm for separable, piecewise-linear concave utilities. For submodular utilities, Garg, Kulkarni, and Kulkarni [18] proposed an $O(n \log n)$ -approximation algorithm, and Li and Vondrák [25] recently obtained the first constant-approximation algorithm for submodular utilities. Last but not least, Barman et al. [8] and Chaudhury, Garg, and Mehta [12] independently developed $O(n)$ -approximation algorithms for the even more general subadditive utilities.

Azar, Buchbinder, and Jain [4] studied an online resource allocation problem whose competitive algorithms are also competitive w.r.t. Nash welfare. However, they assumed that any agent's values for different items could only differ by a bounded ratio; their competitive ratio is logarithmic in the numbers of agents and items as well as this bounded ratio. By contrast, our notions of balance and impartiality compare different agents' utilities. Their assumption rules out the possibility that an agent may have zero value for some item. Therefore, their results do not apply to our model. Besides that, there is a significant line of work about online resource allocation, which focuses on different objectives, e.g., Gkatzelis et al. [19] considered maximizing utilitarian social welfare with envy-freeness in online resource allocation. Banerjee et al. [6] studied an online public items allocation problem with predictions, where the objective is to achieve proportional fairness.

2 Preliminaries

2.1 Nash Welfare Maximization

Consider N agents and T divisible items. Denote the supply of an item t by s_t . We represent an allocation of items to agents by a non-negative real vector $\mathbf{x} = (x_{it})_{1 \leq i \leq N, 1 \leq t \leq T}$, where x_{it} is the amount of item t allocated to agent i . A feasible allocation must comply with the supply constraints of items, that is, it must satisfy $\sum_{i=1}^N x_{it} \leq s_t$ for any item t . For each agent $1 \leq i \leq N$ and each item $1 \leq t \leq T$, $v_{it} \geq 0$ denotes agent i 's value for receiving a unit of item t . The agents' utilities are *additive*: If an agent i receives x_{it} amount of each item t , its utility equals $u_i = \sum_{t=1}^T x_{it}v_{it}$.

Definition 1. The *Nash welfare* of an allocation $\mathbf{x} = (x_{it})_{1 \leq i \leq N, 1 \leq t \leq T}$ is the geometric mean of the agents' utilities:

$$\left(\prod_{i=1}^N u_i \right)^{1/N} = \prod_{i=1}^N \left(\sum_{t=1}^T x_{it}v_{it} \right)^{1/N} .$$

We also consider an equivalent objective, namely, maximizing the sum of logarithms of utilities:

$$\max \sum_{i=1}^N \log u_i .$$

Let $\mathbf{x}^* = (x_{it}^*)_{1 \leq i \leq N, 1 \leq t \leq T}$ denote the allocation that maximizes the Nash welfare, breaking ties arbitrarily. Let $\mathbf{u}^* = (u_i^*)_{1 \leq i \leq N}$ denote the corresponding utilities of the agents.

Maximizing the Nash welfare is a good proxy for balancing the efficiency and fairness of the allocation. Concretely, when the items are divisible, it implies that any agent would get at least $\frac{1}{N}$ of its monopolist utility for receiving all items, a notion of fairness known as proportionality.

Lemma 1 (c.f., Vazirani [30]). *The Nash welfare maximizing allocation \mathbf{x}^* and the corresponding utilities \mathbf{u}^* satisfy that for any agent $1 \leq i \leq N$:*

$$u_i^* \geq \frac{1}{N} \sum_{t=1}^T s_t v_{it} .$$

Proof. Suppose for contrary that there is an agent i such that $u_i^* < \frac{1}{N} \sum_{t=1}^T s_t v_{it}$. We argue that it would be beneficial to allocate more to agent i . Concretely, consider an alternative allocation $\mathbf{x}' = (x'_{it})_{1 \leq i \leq N, 1 \leq t \leq T}$ defined as:

$$x'_{jt} = \begin{cases} \frac{x_{jt}^*}{1 + \varepsilon} & , \text{ if } j \neq i; \\ \frac{x_{it}^* + \varepsilon s_t}{1 + \varepsilon} & , \text{ if } j = i. \end{cases}$$

This is a feasible allocation because for any item $1 \leq t \leq T$:

$$\sum_{j=1}^N x'_{jt} = \frac{1}{1 + \varepsilon} \left(\sum_{j \neq i} x_{jt}^* + (x_{it}^* + \varepsilon s_t) \right) \leq \frac{1}{1 + \varepsilon} (s_t + s_t \varepsilon) = s_t .$$

Further comparing the agents' utilities to the counterparts in the assumed optimal allocation, we have:

$$u_i' = \sum_{t=1}^T x'_{it} v_{it} = \frac{1}{1 + \varepsilon} \left(\sum_{t=1}^T x_{it}^* v_{it} + \varepsilon \sum_{t=1}^T s_t v_{it} \right) = \frac{1}{1 + \varepsilon} \left(u_i^* + \varepsilon \sum_{t=1}^T s_t v_{it} \right) ,$$

and for any agent $j \neq i$ we have $u'_j = \frac{1}{1+\varepsilon}u_j^*$. Hence, we get a contradiction that:

$$\frac{\left(\prod_{i=1}^N u'_i\right)^{1/N}}{\left(\prod_{i=1}^N u_i^*\right)^{1/N}} = \frac{\left(1 + \varepsilon \frac{\sum_{t=1}^T s_t v_{it}}{u_i^*}\right)^{1/N}}{1 + \varepsilon} > 1$$

for a sufficiently small ε by the assumption of $u_i^* < \frac{1}{N} \sum_{t=1}^T s_t v_{it}$. \square

2.2 Online Algorithms

This paper considers an online setting in which the items arrive one by one, while the agents are known from the beginning. When each item arrives, the algorithm observes the item's supply and the agents' values for the item. It must then decide how to allocate the item immediately and irrevocably. We consider the ratio of the optimal Nash welfare to the expected Nash welfare of the algorithm's allocation. The *competitive ratio* of an algorithm is the maximum of this ratio over all possible instances.

Greedy Algorithms with Anticipated Utilities. A natural strategy to allocate an item t is to first estimate how much utility each agent i could get from the other items, denoted as u'_{it} and referred to as agent i 's anticipated utility. We may then allocate item t greedily to maximize the Nash welfare based on the anticipated utilities. For example, a natural yet conservative form of anticipated utilities in our model would be the agents' utilities from the previous items. In other models, the anticipated utilities may depend on some prior knowledge of the agents' values; for instance, the algorithm of Banerjee et al. [5] may be viewed as a greedy algorithm with anticipated utilities that depend on the agents' monopolist utilities.

It is therefore instructive to first examine some basic properties of this family of algorithms. We may formulate the allocation problem conditioned on anticipated utilities u'_{it} as a convex program:

$$\begin{aligned} & \text{maximize} && \sum_{i=1}^N \log(u'_{it} + v_{it}x_{it}) && (1) \\ & \text{subject to} && \sum_{i=1}^N x_{it} \leq s_t \quad \text{and} \quad x_{it} \geq 0 \text{ for any agent } i. \end{aligned}$$

Since the objective is monotone in x_{it} 's, the optimal solution to the above program, denoted as $\mathbf{z}_t = (z_{it})_{1 \leq i \leq N}$, must satisfy the first constraint with equality, i.e.:

$$\sum_{i=1}^N z_{it} = s_t . \quad (2)$$

Further, for the optimal multiplier $\nu^* \geq 0$ of the first constraint and for any agent $1 \leq i \leq N$, the optimality conditions imply that:

$$z_{it} = \max \left\{ \frac{1}{\nu^*} - \frac{u'_{it}}{v_{it}}, 0 \right\} . \quad (3)$$

We give the following lemma to bound the gain of the logarithm of Nash welfare when we follow the optimal solution from Eqn. (3) to allocate an item t .

Lemma 2. *The optimal solution \mathbf{z} to convex program (1) satisfies that:*

$$\sum_{i=1}^N \log(u'_{it} + v_{it}z_{it}) - \sum_{i=1}^N \log u'_{it} \geq s_t \cdot \max_{1 \leq i \leq N} \frac{v_{it}}{u'_{it} + v_{it}z_{it}} .$$

Proof. By the concavity of logarithm, the left-hand side is at least:

$$\sum_{i=1}^N \frac{v_{it}z_{it}}{u'_{it} + v_{it}z_{it}}$$

Further by Eqn. (3), if $z_{it} > 0$ then i maximizes $\frac{v_{it}}{u'_{it} + v_{it}z_{it}}$ (with maximum value ν^*). Hence, the above equals:

$$\max_{1 \leq i \leq N} \frac{v_{it}}{u'_{it} + v_{it}z_{it}} \sum_{i=1}^N z_{it} .$$

The lemma now follows by Eqn. (2). □

Randomized versus Deterministic Algorithms. For divisible items, the best competitive ratios achievable by randomized and deterministic online algorithms are the same. Given any randomized algorithm, we may convert it into a deterministic one such that for each item t , the amount of item t that the deterministic algorithm allocates to any agent i equals the expected amount that the randomized algorithm would allocate. Since the Nash welfare is a concave function of the allocation, the deterministic algorithm yields a weakly larger Nash welfare. Therefore, our positive results will describe online algorithms in their randomized form if that simplifies the presentation. In our negative results, on the other hand, it suffices to consider deterministic algorithms.

2.3 Balanced and Impartial Instances

Banerjee et al. [5] showed that for arbitrary instances, no online algorithm can have a competitive ratio better than the trivial $O(N)$. They then explored the model of online algorithms with predictions, where the algorithms know the agents' monopolist utilities for receiving all items.

We observe that the agents' values in the hard instances of Banerjee et al. [5] are lopsided: the largest value is exponentially larger than the smallest. This is rarely the case in real resource allocation problems. Recall the example of allocating food and computational resources. It is almost impossible to see agents' values towards the same item differ substantially. Hence this paper aims to design algorithms with good performance on "natural instances" in which different agents are of "similar importance". We define two notions of "natural instances". Motivated by Banerjee et al. [5], our first notion compares the agents' monopolist utilities for receiving all items.

Definition 2. An instance is λ -balanced if for any agents $1 \leq i, j \leq N$:

$$\sum_{t=1}^T s_t v_{it} \leq \lambda \sum_{t=1}^T s_t v_{jt} .$$

We shall refer to $\lambda^* = \frac{\max_{1 \leq i \leq N} \sum_{t=1}^T s_t v_{it}}{\min_{1 \leq i \leq N} \sum_{t=1}^T s_t v_{it}}$ as the *balance ratio* of the instance.

We remark again that competitive online algorithms for λ -balanced instances can be transformed into online algorithms with predictions in the model of Banerjee et al. [5]. Given prediction \tilde{V}_i of

each agent i 's monopolist utility $V_i = \sum_{t=1}^T s_t v_{it}$ that is (α, β) -approximate, i.e., $\alpha V_i \geq \tilde{V}_i \geq \frac{1}{\beta} V_i$, we can normalize agent i 's values to be $v'_{it} = \frac{v_{it}}{V_i}$ for all items t to get an $\alpha\beta$ -balanced instance.

Our second notion examines whether different agents get similar utilities in the Nash welfare maximizing allocation.

Definition 3. Given any instance, an allocation \mathbf{x} is μ -impartial if the corresponding utilities \mathbf{u} satisfy that for any agents $1 \leq i, j \leq N$:

$$u_i \leq \mu \cdot u_j .$$

An instance is μ -impartial if every Nash welfare maximizing allocation \mathbf{x}^* is μ -impartial. Let μ^* be the smallest value of μ such that the instance is μ -impartial; we shall refer to μ^* as the *impartiality ratio* of the instance.

Impartiality has a simple yet important implication: For each item, we may ignore the agents whose values for the item are much smaller than the highest value. We make this precise with the next lemma.

Lemma 3. For any agent i and any item t , if $x_{it}^* > 0$ then $v_{it} \geq \frac{1}{\mu^*} \max_{1 \leq j \leq N} v_{jt}$.

Proof. Suppose for contradiction that there are agent i and item t for which we have $x_{it}^* > 0$ but $v_{it} < \frac{1}{\mu^*} \max_{1 \leq j \leq N} v_{jt}$. We argue that we should have allocated less item t to agent i . Let j be an agent with highest value v_{jt} for item t . Consider reallocating an ε amount of item t from agent i to agent j . We claim that the Nash welfare increases for a sufficiently small ε . Since the other agents' utilities stay the same, it suffices to prove that:

$$(u_i^* - \varepsilon v_{it})(u_j^* + \varepsilon v_{jt}) - u_i^* u_j^* = \varepsilon(u_i^* v_{jt} - u_j^* v_{it}) - \varepsilon^2 v_{it} v_{jt}$$

is positive. Since $v_{it} < \frac{1}{\mu^*} v_{jt}$ by our assumption for contradiction and $u_j^* \leq \mu^* u_i^*$ by the definition of μ^* , we get that $u_i^* v_{jt} - u_j^* v_{it} > 0$. Hence, the above is positive for a sufficiently small ε . \square

Finally, the balance ratio and impartiality ratio are within a factor N from each other.

Lemma 4. For any instance, we have $\lambda^* \leq \mu^* N$ and $\mu^* \leq \lambda^* N$.

Proof. We first prove that $\lambda^* \leq \mu^* N$. Let i and j be the agents with the maximum and minimum values for receiving all items respectively. We have:

$$\begin{aligned} \sum_{t=1}^T s_t v_{it} &= \lambda^* \sum_{t=1}^T s_t v_{jt} \geq \lambda^* u_j^* \\ &\geq \frac{\lambda^*}{\mu^*} u_i^* && \text{(definition of } \mu^*) \\ &\geq \frac{\lambda^*}{\mu^* N} \sum_{t=1}^T s_t v_{it} . && \text{(Lemma 1)} \end{aligned}$$

Cancelling $\sum_{t=1}^T s_t v_{it}$ on both sides proves the inequality.

Next, we show that $\mu^* \leq \lambda^* N$ using a similar argument. Let i and j be the agents with the minimum and maximum utilities in the Nash welfare maximizing allocation respectively. We have:

$$\begin{aligned} \sum_{t=1}^T s_t v_{it} &\geq \frac{1}{\lambda^*} \sum_{t=1}^T s_t v_{jt} && \text{(definition of } \lambda^*) \\ &\geq \frac{1}{\lambda^*} u_j^* = \frac{\mu^*}{\lambda^*} u_i^* \\ &\geq \frac{\mu^*}{\lambda^* N} \sum_{t=1}^T s_t v_{it} . && \text{(Lemma 1)} \end{aligned}$$

Cancelling $\sum_{t=1}^T s_t v_{it}$ on both sides proves the inequality. \square

In conclusion, these two notions reflect the “naturalness” of an instance, as a replacement of prediction in previous works. In Section 3, we study balanced instances and achieve $O(\log \lambda^* N)$ -competitive ratio. For impartial instances, besides the $O(\log \mu^* N)$ result as a corollary of Lemma 4, in Section 4 we give an online algorithm whose competitive ratio is $O(\log^2 \mu^*)$, only depending on the impartiality ratio. (All log log factors are omitted here.) Therefore, hopefully the impartiality ratio is the better one to use.

3 Balanced Instances

This section studies balanced instances. Section 3.1 introduces our algorithm under an additional assumption that we were given an upper bound of the balance ratio (Section 3.1). Then, Section 3.2 explains how to remove this assumption by guessing the balance ratio, while losing at most log log factors in the competitive ratio.

3.1 Algorithm with a Known Upper Bound of the Balance Ratio

Suppose that we are given an upper bound λ of the instance’s balance ratio λ^* . In other words, we know that the instance is λ -balanced. The algorithm will divide each item t ’s supply equally into two halves. On the one hand, it allocates the first half equally to all agents, a naïve strategy that is certainly fair but is not efficient enough to approximately maximize the Nash welfare on its own. On the other hand, it greedily allocates the second half of the item to maximize the Nash welfare assuming that each agent would get not only their utilities for the second halves of the previous items allocated to them, but also a fraction of the sum of *all agents*’ monopolist utilities for all known items. Concretely, let $\mathbf{z}_t = (z_{it})_{1 \leq i \leq N}$ denote the allocation of the second half of item t . For anticipated utilities:

$$u'_{it} = \frac{1}{2\lambda N^2} \underbrace{\sum_{j=1}^N \sum_{t'=1}^t v_{jt'} s_{t'}}_{\text{sum of monopolist utilities for all known items, including item } t} + \underbrace{\sum_{t'=1}^{t-1} v_{it'} z_{it'}}_{i\text{'s utility for the second halves of previous items allocated to } i, \text{ excluding item } t} \quad (4)$$

the algorithm chooses \mathbf{z} to maximize $\sum_{i=1}^N \log(u'_{it} + v_{it} x_{it})$ subject to $\sum_{i=1}^N z_{it} \leq \frac{s_t}{2}$ and $z_{it} \geq 0$ for all agents i . We call this algorithm Half-and-Half. See Algorithm 1 for a formal definition.

Theorem 5. *Algorithm 1 is $O(\log \lambda N)$ -competitive.*

Algorithm 1: Half-and-Half (for λ -balanced instances)

- 1 **for** each item $1 \leq t \leq T$ **do**
 - 2 Let $y_{it} = \frac{s_t}{2N}$ for all agents $1 \leq i \leq N$.
 - 3 Let z_{it} maximize (for anticipated utilities u'_{it} defined in Eqn. (4))

$$\sum_{i=1}^N \log(u'_{it} + v_{it}z_{it})$$
 - 4 subject to $\sum_{i=1}^N z_{it} \leq \frac{s_t}{2}$ and $z_{it} \geq 0$ for all agents $1 \leq i \leq N$.
 - 4 Allocate $x_{it} = y_{it} + z_{it}$ amount of item t to each agent $1 \leq i \leq N$.
-

Before getting into the proof of Theorem 5, a comparison with the Set-Aside Greedy algorithm of Banerjee et al. [5] is warranted, since readers familiar with the previous algorithm may have noticed the similarity between the two algorithms. Both algorithms divide each item equally into two halves; both allocate the first half equally, and the second half by some greedy algorithm with anticipated utilities. The difference lies in the designs of anticipated utilities. The Set-Aside Greedy algorithm may be viewed as replacing the first part of our anticipated utility in Eqn. (4) by a $\frac{1}{2N}$ fraction of the prediction on agent i 's monopolist utility. Following the idea of Set-Aside Greedy, a natural attempt is to use each agent i 's monopolist utility for all known items as the prediction of its final monopolist utility. In the online setting, however, the items that contribute the most to an agent's monopolist utility may come at the very end. In that case, the algorithm would underestimate the agent's final utility for a long time, and as a result might unnecessarily allocate many items to this agent in the early rounds. Our solution is to aggregate the monopolist utilities of *all agents* for the known items into an anticipated utility for *every agent*, an idea driven by the assumption of balanced instances.

We next present the analysis of Algorithm 1. It is useful to define the following auxiliary utilities for any agent $1 \leq i \leq N$ and any item $1 \leq t \leq T$:

$$\hat{u}_{it} = \frac{1}{2\lambda N^2} \underbrace{\sum_{j=1}^N \sum_{t'=1}^t v_{jt'} s_{t'}}_{\text{sum of monopolist utilities for items 1 to } t} + \underbrace{\sum_{t'=1}^t v_{it'} z_{it'}}_{i\text{'s utility for the second halves of items 1 to } t \text{ allocated to } i} .$$

By the underlying logic of the algorithm's greedy allocation of the second halves of the items, this is what the algorithm anticipates agent i 's utility to be after allocating item t . The next lemma validates this anticipation at the end of the algorithm.

Lemma 6. *For each agent i , $u_i \geq \hat{u}_{iT}$.*

Proof. It suffices to show that agent i 's utility for the first halves of the items allocated to it is greater than or equal to the first part of \hat{u}_{iT} . By definition, agent i 's utility for the first halves is:

$$\frac{1}{2N} \sum_{t=1}^T v_{it} s_t .$$

Further by the assumption of balanced instances, the monopolist utility of any other agent is

at most λ times larger than agent i 's monopolist utility. Therefore, we have:

$$\sum_{j=1}^N \sum_{t=1}^T v_{jt} s_t \leq \lambda N \sum_{t=1}^T v_{it} s_t .$$

Combining the two claims proves the lemma. \square

Proof of Theorem 5. Since the second halves of the items are allocated by a greedy algorithm with anticipated utilities, we may use Lemma 2 in the analysis. For Lemma 2 to be effective, however, we need the anticipated utilities to be good approximations of the agents' utilities in the Nash welfare maximizing allocation. Fortunately, we only need a polynomial approximation, which is satisfied sufficiently early so that allocating the remaining items correctly still yields approximately optimal Nash welfare.

Let t^* be the earliest item for which $\sum_{i=1}^N \sum_{t=1}^{t^*} v_{it} \geq \frac{1}{2} \min_{1 \leq i \leq N} u_i^*$. The choice of t^* ensures two properties. First, the contribution of items from t^* to T to any agent i 's utility in the optimal solution is at least $\frac{1}{2} u_i^*$. In other words, even if we had made completely wrong allocations, resulting in zero utilities for all agents, we would have lost at worst half of the Nash welfare. Second, the anticipated utility is from now on at least a polynomial approximation of the minimum utility of an agent in the Nash welfare maximizing allocation. Further, the instance is at worst $N\lambda$ -impartial according to Lemma 4. Hence, this is in fact a polynomial approximation for every agent.

When an item $t \geq t^*$ arrives, by Lemma 2 the allocation $\mathbf{z}_t = (z_{it})_{1 \leq i \leq N}$ of the second half of this item satisfies:

$$\sum_{i=1}^N \left(\log(u'_{it} + v_{it} z_{it}) - \log u'_{it} \right) \geq \frac{s_t}{2} \max_{1 \leq i \leq N} \frac{v_{it}}{u'_{it} + v_{it} z_{it}} . \quad (5)$$

To relate our inequality to the Nash welfare maximizing allocation \mathbf{x}^* , we apply $\sum_{i=1}^N x_{it}^* \leq s_t$ to lower bound the right-hand side of Eqn. (5) by:

$$\frac{1}{2} \sum_{i=1}^N \frac{v_{it} x_{it}^*}{u'_{it} + v_{it} z_{it}} .$$

Further, by Lemma 6 we have $u_i \geq \hat{u}_{iT}$ for any agent i , while \hat{u}_{iT} is greater than or equal to \hat{u}_{it} by the definition of \hat{u}_{it} 's. Putting together, we conclude that:

$$\sum_{i=1}^N \left(\log(u'_{it} + v_{it} z_{it}) - \log u'_{it} \right) \geq \frac{1}{2} \sum_{i=1}^N \frac{v_{it} x_{it}^*}{u_i} .$$

Since $\hat{u}_{it} = u'_{it} + v_{it} z_{it}$ and $\hat{u}_{i(t-1)} \leq u'_{it}$, we get telescopic cancellations summing over items from t^* to T . Note that, however, we shall not apply this relaxation for u'_{it^*} for a technical reason that shall be clear shortly. On the other hand, the numerator on the right-hand side sums to at least $\frac{u_i^*}{2}$ for all agents i because of the choice of t^* . Hence:

$$\sum_{i=1}^N \left(\log \hat{u}_{iT} - \log u'_{it^*} \right) \geq \frac{1}{4} \sum_{i=1}^N \frac{u_i^*}{u_i} .$$

Algorithm 2: Half-and-Half (for instances with unknown balance ratio)

- 1 Sample λ such that it equals 2^{2^k} with probability $\frac{6}{\pi^2} \cdot \frac{1}{(k+1)^2}$ for any non-negative integer k .
 - 2 Run Algorithm 1 with λ to allocate the items to the agents.
-

Further, by Lemma 6 we have $u_i \geq \hat{u}_{iT}$. We also have:

$$\begin{aligned}
u'_{it^*} &\geq \frac{1}{2\lambda N^2} \sum_{i=1}^N \sum_{t=1}^{t^*} v_{it} && \text{(definition of } u'_{it}\text{)} \\
&\geq \frac{1}{4\lambda N^2} \min_{1 \leq i \leq N} u_i^* && \text{(choice of } t^*\text{)} \\
&\geq \frac{1}{4\lambda^2 N^3} \left(\prod_{i=1}^N u_i^* \right)^{\frac{1}{n}}. && \text{(\lambda N-impartiality by Lemma 4)}
\end{aligned}$$

Note that this inequality would not be true in general if we had relaxed u'_{it^*} to $\hat{u}_{i(t^*-1)}$. Putting together and by AM-GM inequality, we get that:

$$\log \prod_{i=1}^N u_i - \log \prod_{i=1}^N u_i^* + N \log 4\lambda^2 N^3 \geq \frac{1}{4} \sum_{i=1}^N \frac{u_i^*}{u_i} \geq \frac{N}{4} \left(\prod_{i=1}^N \frac{u_i^*}{u_i} \right)^{\frac{1}{N}}.$$

Let $\Gamma = \left(\prod_{i=1}^N \frac{u_i^*}{u_i} \right)^{\frac{1}{N}}$ be the ratio of the optimal Nash welfare to the algorithm's Nash welfare. The above inequality is equivalent to:

$$-\log \Gamma + \log 4\lambda^2 N^3 \geq \frac{1}{4} \Gamma.$$

Since $\Gamma \geq 1$ and thus $\log \Gamma \geq 0$, the above inequality implies $\Gamma \leq 4 \log 4\lambda^2 N^3 = O(\log \lambda N)$. \square

3.2 Guessing the Balance Ratio

When we have no prior knowledge of the balance ratio, we can guess an upper bound of the balance ratio by sampling from an appropriate distribution. Since the final ratio depends logarithmically on the upper bound, it suffices to make a good enough guess that is at most a polynomial of the true balance ratio λ^* .

Concretely, we shall consider a sequence of numbers starting from 2, such that each sequel number is the square of the previous number. We will sample each number λ with a probability that is inverse polynomial in $\log \log \lambda$; this ensures that the correct guess is made with a sufficiently large probability. Finally, we apply the prior-dependent Half-and-Half algorithm (Algorithm 1) with the guessed upper bound λ . See Algorithm 2 for a formal definition.

Theorem 7. *Algorithm 2 is $O(\log \lambda^* N (\log \log \lambda^*)^2)$ -competitive.*

Proof. Suppose that k is the smallest positive integer such that the balance ratio λ^* is at most 2^{2^k} . Then, we have $2^{2^{k-1}} < \lambda^* \leq 2^{2^k}$ which further implies that $\log \log \lambda^* > k - 1$. Hence, the algorithm correctly guesses $\lambda = 2^{2^k}$ with probability at least $\Omega(\frac{1}{k^2}) = \Omega(\frac{1}{(\log \log \lambda^*)^2})$. When that happens, the algorithm's Nash welfare is an $O(\log \lambda^* N)$ approximation to the optimal Nash welfare. Therefore, even if the algorithm got zero Nash welfare from the other guesses, we would still have the stated competitive ratio. \square

Considering the relation of balance and impartiality ratios (Lemma 4), Theorem 7 can directly imply Theorem 8, which means the algorithm is also competitive for impartial instances. The next section will develop algorithms tailored for impartial instances with competitive ratios independent of the number of agents.

Theorem 8. *Algorithm 2 is $O(\log \mu^* N (\log \log \mu^* N)^2)$ -competitive.*

4 Impartial Instances

This section considers impartial instances. Section 4.1 examines a special case of binary values: an agent i 's value for an item t is either $v_{it} = 0$, or some value $v_{it} = v_t$ that is the same for all agents. We will show that a simple greedy algorithm, which was referred to as Myopic Greedy by Banerjee et al. [5], is $O(\log \mu^*)$ -competitive in this special case. This is nearly tight as the next section will show an almost matching lower bound. Section 4.2 then explains how to reduce the general case to the binary case, under an additional assumption that we are given an upper bound μ of the impartiality ratio μ^* , i.e., if we know that the instance is μ -impartial. Finally, Section 4.3 applies the same technique as in the previous section to randomly guess an upper bound, removing the additional assumption while losing a factor that depends on $\log \log \mu^*$ in the competitive ratio.

4.1 Myopic Greedy and Binary Values

The Myopic Greedy algorithm simply allocates each item t greedily to maximize the Nash welfare conditioned on the allocation before item t . In other words, it is a greedy algorithm with anticipated utilities, for which the anticipated utility of an agent equals its utility for the items allocated to it so far. See Algorithm 3 for a formal definition.

Algorithm 3: Myopic Greedy

1 **for** each item $1 \leq t \leq T$ **do**
2 Let the allocation $\mathbf{x}_t = (x_{it})_{1 \leq i \leq N}$ maximize:

$$\sum_{i=1}^N \log \left(\sum_{t'=1}^{t-1} v_{it'} x_{it'} + v_{it} x_{it} \right)$$

 subject to $\sum_{i=1}^N x_{it} \leq s_t$ and $x_{it} \geq 0$ for all agents $1 \leq i \leq N$.

The rest of this subsection will assume that the agents' values are binary, that is, for any item $1 \leq t \leq T$ and any agent $1 \leq i \leq N$, either $v_{it} = 0$ or $v_{it} = v_t$.

Theorem 9. *Algorithm 3 is $O(\log \mu^*)$ -competitive if the agents have binary values for the items.*

In fact, we will prove a slightly stronger result so that in the next subsection we can reduce the general case to the case of binary values. We formulate the stronger claim as the next lemma.

Lemma 10. *For any feasible allocation $\tilde{\mathbf{x}} = (\tilde{x}_{it})_{1 \leq i \leq N, 1 \leq t \leq T}$ and the agents' corresponding utilities $\tilde{\mathbf{u}} = (\tilde{u}_i)_{1 \leq i \leq N}$, if allocation $\tilde{\mathbf{x}}$ is $\tilde{\mu}$ -impartial, then the Nash welfare of the allocation by Algorithm 3 is at least an $O(\log \tilde{\mu})$ approximation to the Nash welfare of allocation $\tilde{\mathbf{x}}$, i.e.:*

$$\left(\frac{\prod_{i=1}^N \tilde{u}_i}{\prod_{i=1}^N u_i} \right)^{\frac{1}{N}} = O(\log \tilde{\mu}).$$

Theorem 9 follows as a corollary by letting $\tilde{\mathbf{x}}$ be the Nash welfare maximizing allocation \mathbf{x}^* .

The rest of the subsection focuses on proving Lemma 10. We assume without loss of generality that the agents are sorted by their utilities for the algorithm's allocation, i.e.:

$$u_1 \leq u_2 \leq \dots \leq u_N .$$

We start with the following lemma which links the agents' utilities \mathbf{u} for the algorithm's allocation, and their utilities $\tilde{\mathbf{u}}$ for the benchmark allocation $\tilde{\mathbf{x}}$.

Lemma 11. *For any $1 \leq i \leq N$:*

$$\sum_{j=1}^i \tilde{u}_j \leq \sum_{j=1}^N \min\{u_i, u_j\}. \quad (6)$$

Proof. For any fixed $1 \leq i \leq N$, let S_i be the set of items for which at least one of agents 1 to i have positive values, i.e.:

$$S_i = \left\{ 1 \leq t \leq T : \exists 1 \leq j \leq i, v_{jt} = v_t \right\} .$$

Then, the left-hand side of Equation (6) is upper bounded by:

$$\sum_{j=1}^i \tilde{u}_j \leq \sum_{t \in S_i} s_t v_t .$$

To prove the lemma, it remains to show that the items in S_i contribute at most $\min\{u_i, u_j\}$ to any agent j 's utility for the algorithm's allocation. For agents $1 \leq j \leq i$, this is trivially true because the stated utility bound is simply u_j . For any agent $j > i$, consider the last moment when it receives a positive amount of some item t in S_i . It suffices to show that its utility is at most u_i after the allocation of item t . Recall that the Myopic Greedy algorithm is a greedy algorithm with anticipated utilities. First by the assumption of binary values and by Equation (3), we conclude that any agent who receives a positive amount of item t must have the smallest utility among all agents with values v_t for item t , after the allocation of item t . Meanwhile, one of the agents from 1 to i has value v_t for the item since $t \in S_i$, and this agent has utility at most u_i . Therefore, agent j 's utility after the allocation of item t is at most u_i as desired. \square

If we view the agents' utilities \mathbf{u} for the algorithm's allocation as variables, Lemma 11 offers a set of linear inequalities that relates them with the benchmark utilities $\tilde{\mathbf{u}}$. The next lemma shows that subject to these inequalities, the smallest Nash welfare w.r.t. \mathbf{u} is achieved when the inequalities all hold with equalities.

Lemma 12. *For any non-negative $\mathbf{u} = (u_i)_{1 \leq i \leq N}$ and $\tilde{\mathbf{u}} = (\tilde{u}_i)_{1 \leq i \leq N}$ that satisfy the inequality in Equation (6), we have:*

$$\sum_{i=1}^N \log u_i \geq \sum_{i=1}^N \log \left(\sum_{j=1}^i \frac{\tilde{u}_i}{N+1-j} \right) . \quad (7)$$

Proof. For ease of exposition, we rewrite Equation (6) as:

$$\begin{cases} u_1 \cdot N & \geq \tilde{u}_1 \\ u_1 \cdot N + (u_2 - u_1) \cdot (N - 1) & \geq \tilde{u}_1 + \tilde{u}_2 \\ \dots & \\ u_1 \cdot N + \sum_{i=1}^{N-1} ((u_{i+1} - u_i) \cdot (N - i)) & \geq \sum_i^N \tilde{u}_i \end{cases} \quad (8)$$

Along with the restrictions $u_1 \leq u_2 \leq \dots \leq u_N$, the feasible region of $(u_i)_{i \in [N]}$ is a polytope. Since $\sum_{i=1}^N \log u_i$ is concave, it reaches minimum at a vertex of the polytope, i.e., where exactly n of the inequalities are equalities. Let $(\bar{u}_i)_{i \in [N]}$ be a solution with minimum $\sum_{i=1}^N \log u_i$. Suppose there exists i such that $\bar{u}_i = \bar{u}_{i+1}$. Let i^* be the smallest such i . The i^* -th inequality in Equation (8) is not tight.

Define $(u'_1, u'_2, \dots, u'_N)$ as

$$\begin{cases} u'_i = \bar{u}_i & \forall i < i^* \\ u'_i = \bar{u}_i - \varepsilon & i = i^* \\ u'_i = \bar{u}_i + \frac{\varepsilon}{N-i} & \forall i > i^* \end{cases} \quad (9)$$

for sufficiently small $\varepsilon > 0$. $(u'_1, u'_2, \dots, u'_N)$ satisfies Equation (8), and $u'_1 u'_2 \dots u'_N < \bar{u}_1 \bar{u}_2 \dots \bar{u}_N$ holds since \bar{u}_i is the smallest among $\bar{u}_i, \bar{u}_{i+1}, \dots, \bar{u}_N$, contradicting the minimality of $(\bar{u}_i)_{i \in [N]}$.

Therefore, $u_i = \sum_{i'=1}^i \frac{\bar{u}_i}{N+1-i'}$ in the solution that minimizes $\sum_{i=1}^N \log u_i$. \square

We will relax the denominators on the right-hand side of Equation (7) to be N in the subsequent analysis. We get that:

$$\sum_{i=1}^N \log u_i \geq \sum_{i=1}^N \log \left(\frac{1}{i} \sum_{j=1}^i \tilde{u}_j \right) - \log \frac{N^N}{N!} . \quad (10)$$

If the benchmark allocation $\tilde{\mathbf{x}}$ was 1-impartial, we would have $\tilde{u}_i = \left(\prod_{j=1}^N \tilde{u}_j \right)^{\frac{1}{N}}$ for all agents $1 \leq i \leq N$. In that case, the above Inequality (10) would already imply a constant competitive ratio because it can be written as:

$$\log \left(\prod_{i=1}^N u_i \right)^{\frac{1}{N}} \geq \log \left(\prod_{i=1}^N \tilde{u}_i \right)^{\frac{1}{N}} - \log \frac{N}{(N!)^{\frac{1}{N}}} ,$$

where the second term on the right-hand side is $O(1)$ by Stirling's approximation.

For μ -impartial instances with an arbitrary $\mu \geq 1$, we need one more inequality given in the next lemma. The form of the inequality is clean so we suspect that it may have already been proved in the past. To our best effort, however, we only found Hardy's Inequality and its several variants (c.f., Hardy, Littlewood, and Pólya [21]) to have a similar spirit as they also compare the function values of a sequence of non-negative real numbers and the functions values of their prefix averages; but those inequalities do not imply our lemma. It would be interesting to find further applications of this inequality in other problems.

Lemma 13. *Suppose that a_1, a_2, \dots, a_N are real numbers between 1 and μ . We have:*

$$\frac{1}{N} \sum_{i=1}^N \log a_i - \frac{1}{N} \sum_{i=1}^N \log \left(\frac{1}{i} \sum_{j=1}^i a_j \right) \leq \log(\log \mu + 1) + 1 .$$

Proof. It is without loss of generality to assume that $a_1 \leq a_2 \leq \dots \leq a_N$ since this order minimizes the second term on the left-hand side conditioned on the set of these N real numbers. We may further assume $a_1 = 1$ because the left-hand side remains the same if a_1, a_2, \dots, a_N are multiplied by the same factor.

Our proof strategy is to consider all possible thresholds $\theta \geq 0$ and to bound the number of indices ($1 \leq i \leq N$) for which $\log a_i - \log \frac{1}{i} \sum_{j=1}^i a_j$ is at least θ . First, we rewrite the left-hand side as:

$$\int_0^\infty \frac{\sum_{i=1}^N \mathbf{1}(\log a_i - \log(\frac{1}{i} \sum_{j=1}^i a_j) \geq \theta)}{N} d\theta ,$$

where $\mathbf{1}(\cdot)$ is the indicator function.

For any $\theta \geq 0$, suppose that $i_1 < i_2 < \dots < i_k$ are the indices for which

$$\log a_i - \log \left(\frac{1}{i} \sum_{j=1}^i a_j \right) \geq \theta .$$

First, by $a_1 = 1 < a_{i_1}$ and $a_j \leq a_{i_1}$ for all $j \leq i$, the left-hand side above is strictly smaller than $\log i_1$ when $i = i_1$. This gives $i_1 > e^\theta$.

For ease of exposition, in the remaining argument we define $i_0 = 1, a_{i_0} = 1$. Then, we have that for any $j \geq 1$:

$$\begin{aligned} \frac{i_j}{e^\theta} a_{i_j} &\geq (a_1 + a_2 + \dots + a_{i_j}) \\ &\geq (i_1 - i_0)a_{i_0} + (i_2 - i_1)a_{i_1} + \dots + (i_j - i_{j-1})a_{i_{j-1}} + a_{i_j} . \end{aligned}$$

Subject to these inequalities, a_{i_k} achieves its minimum value when equalities hold everywhere. In that case, for any $j \geq 1$ we have (by taking the difference of two consecutive equalities):

$$\frac{i_{j+1}}{e^\theta} a_{i_{j+1}} - \frac{i_j}{e^\theta} a_{i_j} = (i_{j+1} - i_j - 1)a_{i_j} + a_{i_{j+1}} .$$

Therefore:

$$\begin{aligned} a_{i_{j+1}} &= \left(e^\theta - \frac{(e^\theta - 1)(\frac{i_j}{e^\theta} - 1)}{\frac{i_{j+1}}{e^\theta} - 1} \right) a_{i_j} \\ &\geq \left(e^\theta - \frac{(e^\theta - 1)(\frac{i_j}{e^\theta} - 1)}{\frac{i_{j+1}}{e^\theta} - 1} \right) a_{i_j} && (i_{j+1} \geq i_j + 1) \\ &\geq \left(e^\theta - \frac{(e^\theta - 1)(\frac{N}{e^\theta} - 1)}{\frac{N+1}{e^\theta} - 1} \right) a_{i_j} && (i_j \leq N) \\ &= \frac{1}{1 - \frac{e^\theta - 1}{N}} a_{i_j} \\ &\geq e^{\frac{\theta - 1}{N}} a_{i_j} . && (1 - y \leq e^{-y}) \end{aligned}$$

For a_{i_1} , we have $a_{i_1} = \frac{e^\theta}{i_1}(i_1 - 1 + a_{i_1})$ which means (recall that $i_1 > e^\theta$):

$$a_{i_1} = \frac{(i_1 - 1)e^\theta}{i_1 - e^\theta} \geq \frac{Ne^\theta}{(N + 1)e^{-\theta} - 1} = \frac{e^\theta}{1 - \frac{e^\theta - 1}{N}} \geq \frac{1}{1 - \frac{e^\theta - 1}{N}} \geq e^{\frac{\theta - 1}{N}} .$$

Combining the above two inequalities, we get that $a_{i_k} \geq e^{\frac{k}{N}(e^\theta - 1)}$. On the other hand, we have $a_{i_k} \leq \mu$. Hence:

$$\frac{k}{N} \leq \frac{\log \mu}{e^\theta - 1} .$$

Algorithm 4: Greedy with Rounded Values (for μ -impartial instances)

```

1 for each item  $1 \leq t \leq T$  do
2   Let  $\bar{v}_t = \max_{1 \leq i \leq N} v_{it}$ 
3   for  $j = 1$  to  $\lceil \log \mu \rceil$  do
4     Let there be a sub-item  $(t, j)$  with supply  $\frac{s_t}{\lceil \log \mu \rceil}$ .
5     Let each agent  $i$ 's value for the sub-item be:
        
$$v_{i(t,j)} = \begin{cases} \frac{\bar{v}_t}{2^j} & \text{if } v_{it} \geq \frac{\bar{v}_t}{2^j} \\ 0 & \text{otherwise.} \end{cases}$$

6     Let Algorithm 3 allocate the sub-items and let the allocation be
         $(x_{i(t,j)})_{1 \leq i \leq N, 1 \leq j \leq \lceil \log \mu \rceil}$ .
7     Allocate  $x_{it} = \sum_{j=1}^{\lceil \log \mu \rceil} x_{i(t,j)}$  amount of item  $t$  to each agent  $1 \leq i \leq N$ .

```

Therefore, the left-hand side of the lemma's inequality is at most:

$$\begin{aligned}
\int_0^\infty \min \left\{ \frac{\log \mu}{e^\theta - 1}, 1 \right\} d\theta &= \int_0^{\log(\log \mu + 1)} 1 d\theta + \int_{\log(\log \mu + 1)}^\infty \frac{\log \mu}{e^\theta - 1} d\theta \\
&= \log(\log \mu + 1) + \log \mu \cdot \log \left(1 + \frac{1}{\log \mu} \right) \\
&\leq \log(\log \mu + 1) + 1,
\end{aligned}$$

where the last inequality follows by $\log(1 + y) \leq y$ for any $y \geq 0$. \square

Lemma 10 then follows by combining Equation (10), Stirling's approximation, and Lemma 13 with $a_i = \frac{\tilde{u}_i}{\min_{1 \leq j \leq N} \tilde{u}_j}$ and $\mu = \tilde{\mu}$.

4.2 Algorithm with a Known Upper Bound of the Impartiality Ratio

Suppose that we are given an upper bound μ of the instance's impartiality ratio μ^* . In other words, we know that the instance is μ -impartial. For any item t , let $\bar{v}_t = \max_{i \in [N]} v_{it}$ denote the maximum value of any agent for item t . By Lemma 3, we should only allocate item t to the agents whose values for item t is at least $\frac{1}{\mu} \bar{v}_t$. A naïve approach is to treat all such agents as if they had value \bar{v}_t for the item, and then apply the Myopic Greedy algorithm; doing so would lose an extra factor μ in the competitive ratio in the worst case. Instead, we divide the item equally into $\lceil \log \mu \rceil$ sub-items, each with supply $\frac{s_t}{\lceil \log \mu \rceil}$. For the j -th sub-item, which we referred to as sub-item (t, j) , the agents' values are rounded down to either $\frac{\bar{v}_t}{2^j}$ or 0. That is, an agent i has value $\frac{\bar{v}_t}{2^j}$ for sub-item (t, j) if its original value for item t is at least as much, and has value 0 for this sub-item otherwise. Then, we run the Myopic Greedy algorithm to allocate the sub-items. See Algorithm 4 for a formal definition.

Theorem 14. *Algorithm 4 is $O(\log^2 \mu)$ -competitive.*

Proof. Recall that $\mathbf{x}^* = (x_{it}^*)_{1 \leq i \leq N, 1 \leq t \leq T}$ denote the Nash welfare maximizing allocation. We will

Algorithm 5: Greedy with Rounded Values (for instances with unknown impartiality ratio)

- 1 Sample μ such that it equals 2^{2^k} with probability $\frac{6}{\pi^2} \cdot \frac{1}{(k+1)^2}$ for any non-negative integer k .
 - 2 Run Algorithm 4 with μ to allocate the items to the agents.
-

consider a corresponding feasible allocation $(x'_{i(t,j)})$ of the sub-items as follows:

$$x'_{i(t,j)} = \begin{cases} \frac{x_{it}^*}{\lceil \log \mu \rceil} & \text{if } \frac{\bar{v}_t}{2^j} \leq v_{it} < \frac{\bar{v}_t}{2^{j-1}} \\ 0 & \text{otherwise.} \end{cases}$$

Suppose that an agent i 's value for item t satisfies $\frac{\bar{v}_t}{2^\ell} \leq v_{it} < \frac{\bar{v}_t}{2^{\ell-1}}$ for integer $\ell \geq 1$. Then its utility for the sub-items (t, j) , $1 \leq j \leq \lceil \log \mu \rceil$, allocated to it would be:

$$\sum_{j=1}^{\lceil \log \mu \rceil} x'_{i(t,j)} v_{i(t,j)} = \sum_{j=\ell}^{\lceil \log \mu \rceil} x'_{i(t,j)} v_{i(t,j)} = \frac{x_{it}^*}{\lceil \log \mu \rceil} \sum_{j=\ell}^{\lceil \log \mu \rceil} \frac{\bar{v}_t}{2^j}.$$

The summation on the right-hand-side is at least $\frac{\bar{v}_t}{2^j}$ and at most $\frac{\bar{v}_t}{2^{j-1}}$. Thus, we conclude that:

$$\frac{x_{it}^* v_{it}}{2 \lceil \log \mu \rceil} \leq \sum_{j=1}^{\lceil \log \mu \rceil} x'_{i(t,j)} v_{i(t,j)} \leq \frac{x_{it}^* v_{it}}{\lceil \log \mu \rceil}.$$

The agents' utilities \mathbf{u}' for the sub-items satisfy that for any agent $1 \leq i \leq N$:

$$\frac{u_i^*}{2 \lceil \log \mu \rceil} \leq u'_i \leq \frac{u_i^*}{\lceil \log \mu \rceil}.$$

This implies two properties. First, the allocation of sub-items is 2μ -impartial since the original allocation \mathbf{x}^* is μ -impartial. Second, the Nash welfare of the allocation of sub-items is at least an $\Omega(\frac{1}{\log \mu})$ fraction of the optimal Nash welfare of the original instance. Therefore, by Lemma 10, we get an $O(\log \mu)$ -competitive allocation comparing to the Nash welfare of sub-item allocation, which implies an $O(\log^2 \mu)$ -competitive allocation w.r.t. the original instance. \square

4.3 Guessing the Impartiality Ratio

This is almost verbatim to the counterpart for balanced instances. When we have no prior knowledge of the impartiality ratio, we can guess an upper bound by sampling from an appropriate distribution. Since the final ratio depends logarithmically on the upper bound, it suffices to make a good enough guess that is at most a polynomial of the true impartiality ratio μ^* .

We shall again consider a sequence of numbers starting from 2, such that each sequel number is the square of the previous number. We will sample each number μ with a probability that is inverse polynomial in $\log \log \mu$; this ensures that the correct guess is made with a sufficiently large probability. Finally, we apply the prior-dependent Greedy with Rounded Values algorithm (Algorithm 4) with the guessed upper bound μ . See Algorithm 5 for a formal definition.

Theorem 15. *Algorithm 5 is $O(\log^2 \mu^* (\log \log \mu^*)^2)$ -competitive.*

Table 2: Illustration of a hard instance. Rows are items and columns are agents. The number in the intersection of the t -th row and the i -th column is agent i 's value for item t .

t	Agent 1	Agent 2	Agent 3	\dots	Agent n
1	n^2	n^2	n^2		n^2
2	0	n^4	n^4		n^4
3	0	0	n^6		n^6
			\dots		
n	0	0	0		n^{2n}

Proof. Suppose that k is the smallest positive integer such that the balance ratio μ^* is at most 2^{2^k} . Then, we have $2^{2^{k-1}} < \mu^* \leq 2^{2^k}$. This further implies that $\log \log \lambda^* > k - 1$. Hence, the algorithm correctly guesses $\mu = 2^{2^k}$ with probability at least $\Omega(\frac{1}{k^2}) = \Omega(\frac{1}{(\log \log \mu^*)^2})$. When that happens, the algorithm's Nash welfare is an $O(\log^2 \mu^* N)$ approximation to the optimal Nash welfare. Therefore, even if the algorithm got zero Nash welfare from the other guesses, we would still have the stated competitive ratio. \square

5 Lower Bounds

We first restate a lower bound by Banerjee et al. [5] under our model.

Lemma 16 (Banerjee et al. [5], Theorem 3). *No online algorithm can achieve a competitive ratio better than $\log^{1-o(1)} N$,² even for 1-balanced instance.*

Next we prove that the logarithmic dependence on the balance ratio or the impartiality ratio is necessary. The main ingredient is the instance given by Table 2, which is a variant of another hard instance by Banerjee et al. [5] (c.f., Theorem 4 therein). The instance has N agents, and $T = N$ items of unit supplies. Agent i 's value for item t is N^{2i} if $t \geq i$, and is 0 otherwise. Recall that for lower bounds it suffices to consider deterministic algorithms. Hence, we may assume without loss of generality that agent i receives the least amount of item $t = i$ among agents i to N in the algorithm's allocation.

Lemma 17. *No algorithm can be better than $\frac{n-1}{e}$ competitive for the instance given by Table 2.*

Proof. For any $1 \leq i \leq n$, by the assumption that agent i receives the least amount of item $t = i$ among agents i to n , we get that:

$$x_{ii} \leq \frac{1}{n - i + 1}.$$

Therefore, agent i 's utility is at most

$$u_i \leq \sum_{t=1}^{i-1} n^{2t} + \frac{n^{2i}}{n - i + 1} \leq \frac{n^{2i}}{n - i + 1} \cdot \frac{n}{n - 1}.$$

On the other hand, if we allocate each item $1 \leq t \leq n$ to agent $i = t$ (which is the optimal allocation), agent i 's utility would be:

$$u_i^* = n^{2i}.$$

²The original theorem by Banerjee et al. [5] only claimed a weaker bound of $\log^{1-\epsilon} N$ but their proof actually showed a slightly stronger bound that we restate here.

Therefore, the competitive ratio of the algorithm is at best:

$$\left(\prod_{i=1}^n \frac{u_i^*}{u_i} \right)^{\frac{1}{n}} \geq \left(\prod_{i=1}^n \frac{n^{2i}}{\frac{n^{2i}}{n-i+1} \cdot \frac{n}{n-1}} \right)^{\frac{1}{n}} = \frac{n-1}{n} (n!)^{\frac{1}{n}} > \frac{n-1}{e}.$$

□

Lemma 18. *The imbalance ratio and impartiality ratio of the instance given by Table 2 are at most n^{2n} .*

We can now derive our lower bounds as corollaries of the above lemmas.

Theorem 19. *No algorithm can be better than $\min \{ \log^{1-o(1)} N \lambda^* , \frac{N-1}{e} \}$ -competitive.*

Proof. If $N \geq \lambda^*$, the stated lower bound is $\log^{1-o(1)} N$ which follows by Lemma 16.

If $N < \lambda^*$ but $N^{2N} \geq \lambda^*$, the stated lower bound is $\log^{1-o(1)} \lambda^*$. Consider the instance given by Table 2 choosing $n = \Theta\left(\frac{\log \lambda^*}{\log \log \lambda^*}\right)$ such that $\lambda^* = n^{2n}$. Note that this implies $N \geq n$. Create $\frac{N}{n}$ copies of this instance so that the number of agents becomes N . Then, the stated lower bound follows by Lemmas 17 and 18.

Finally if $N^{2N} < \lambda^*$. The stated bound is $\frac{N-1}{e}$, which follows by Lemmas 17 and 18, choosing $n = N$ in the instance given by Table 2. □

Theorem 20. *No algorithm can be better than $\min \{ \log^{1-o(1)} \mu^* , \frac{N-1}{e} \}$ -competitive.*

Proof. If $N^{2N} \geq \mu^*$, the stated lower bound is $\log^{1-o(1)} \mu^*$. Consider the instance given by Table 2 choosing $n = \Theta\left(\frac{\log \mu^*}{\log \log \mu^*}\right)$ such that $\mu^* = n^{2n}$. Note that this implies $N \geq n$. Create $\frac{N}{n}$ copies of this instance so that the number of agents becomes N . Then, the stated lower bound follows by Lemmas 17 and 18.

Finally if $N^{2N} < \mu^*$. The stated lower bound is $\frac{N-1}{e}$, which follows by Lemmas 17 and 18, choosing $n = N$ in the instance given by Table 2. □

We remark that Theorem 20 still holds even if the agents' values for the items are either 0 or 1. For each item t in Table 2, instead of letting it have unit supply and letting agents t to n have values n^{2t} for it, we can instead let it have supply n^{2t} and let agents t to n have values 1 for it. Scaling the allocation of item t in the original instance by n^{2t} factor would give an allocation for the 0-1 value version, with the same utilities for all agents.

6 Conclusion

In this paper we initiate the study of online Nash welfare maximization *without predictions*. We define *balance ratio* and *impartiality ratio* for an instance, and design online algorithms whose competitive ratios only depend on the logarithms of the aforementioned ratios of agents' utilities and the number of agents. One possible extension of our work is to close the gap between the lower and upper bounds in general instances. In addition to that, our technique may provide an insight on how to remove predictions for other online problems.

References

- [1] 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.
- [2] Nima Anari, Shayan Oveis Gharan, Amin Saberi, and Mohit Singh. Nash social welfare, matrix permanent, and stable polynomials. In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference*, page 36, 2017.
- [3] Nima Anari, Tung Mai, Shayan Oveis Gharan, and Vijay V Vazirani. Nash social welfare for indivisible items under separable, piecewise-linear concave utilities. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2274–2290, 2018.
- [4] Yossi Azar, Niv Buchbinder, and Kamal Jain. How to allocate goods in an online market? In *Proceedings of the 18th Annual European Conference on Algorithms*, page 51–62, 2010.
- [5] 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, 2022.
- [6] Siddhartha Banerjee, Vasilis Gkatzelis, Safwan Hossain, Billy Jin, Evi Micha, and Nisarg Shah. Proportionally fair online allocation of public goods with predictions. *arXiv preprint arXiv:2209.15305*, 2022.
- [7] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 557–574, 2018.
- [8] 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*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- [9] Siddharth Barman, Arindam Khan, and Arnab Maiti. Universal and tight online algorithms for generalized-mean welfare. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 4793–4800, 2022.
- [10] Niv Buchbinder and Joseph Naor. Online primal-dual algorithms for covering and packing. *Mathematics of Operations Research*, 34(2):270–286, 2009.
- [11] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3):1–32, 2019.
- [12] Bhaskar Ray Chaudhury, Jugal Garg, and Ruta Mehta. Fair and efficient allocations under subadditive valuations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5269–5276, 2021.
- [13] Richard Cole and Vasilis Gkatzelis. Approximating the Nash social welfare with indivisible items. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*, pages 371–380, 2015.

- [14] 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 18th ACM Conference on Economics and Computation*, pages 459–460, 2017.
- [15] 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.
- [16] Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.
- [17] Jugal Garg, Martin Hoefer, and Kurt Mehlhorn. Approximating the Nash social welfare with budget-additive valuations. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2326–2340, 2018.
- [18] 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.
- [19] Vasilis Gkatzelis, Alexandros Psomas, and Xizhi Tan. Fair and efficient online allocations with normalized valuations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5440–5447, 2021.
- [20] Fang Hao, Murali Kodialam, TV Lakshman, and Sarit Mukherjee. Online allocation of virtual machines in a distributed cloud. *IEEE/ACM Transactions on Networking*, 25(1):238–249, 2016.
- [21] Godfrey H Hardy, John E Littlewood, and George Pólya. *Inequalities*. Cambridge University Press, 1952.
- [22] Mamoru Kaneko and Kenjiro Nakamura. The Nash social welfare function. *Econometrica*, 47(2):423–435, 1979.
- [23] Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. An optimal algorithm for online bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*, pages 352–358, 1990.
- [24] Euiwoong Lee. APX-hardness of maximizing Nash social welfare with indivisible items. *Information Processing Letters*, 122:17–20, 2017.
- [25] Wenzheng Li and Jan Vondrák. A constant-factor approximation algorithm for nash social welfare with submodular valuations. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science*, pages 25–36, 2022.
- [26] Aranyak Mehta. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science*, 8(4):265–368, 2013.
- [27] Hervé Moulin. *Fair division and collective welfare*. MIT Press, 2004.
- [28] John F Nash Jr. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.
- [29] Canice Prendergast. How food banks use markets to feed the poor. *Journal of Economic Perspectives*, 31(4):145–162, 2017.

- [30] Vijay V Vazirani. Combinatorial algorithms for market equilibria. In *Algorithmic Game Theory*, pages 103–134. Cambridge University Press, 2007.