Greedy-Based Online Fair Allocation with Adversarial Input: Enabling Best-of-Many-Worlds Guarantees

Zongjun Yang[1], **Luofeng Liao**[2], **Christian Kroer**[2*]

[1]School of Electronics Engineering and Computer Science, Peking University
[2]Columbia University
allenyzj@stu.pku.edu.cn, {ll3530, christian.kroer}@columbia.edu

Abstract

We study an online allocation problem with sequentially arriving items and adversarially chosen agent values, with the goal of balancing fairness and efficiency. Our goal is to study the performance of algorithms that achieve strong guarantees under other input models such as stochastic inputs, in order to achieve robust guarantees against a variety of inputs. To that end, we study the PACE (Pacing According to Current Estimated utility) algorithm, an existing algorithm designed for stochastic input. We show that in the equal-budgets case, PACE is equivalent to an integral greedy algorithm. We go on to show that with natural restrictions on the adversarial input model, both the greedy allocation and PACE have asymptotically bounded multiplicative envy as well as competitive ratio for Nash welfare, with the multiplicative factors either constant or with optimal order dependence on the number of agents. This completes a "best-of-many-worlds" guarantee for PACE, since past work showed that PACE achieves guarantees for stationary and stochastic-but-non-stationary input models.

1 Introduction

We study *online fair allocation*, where items arrive sequentially in T rounds, and we need to distribute them among a set of n agents with heterogeneous preferences, with the goal of achieving both fairness and efficiency properties. In each round, we observe each agent's value for the item and make an irrevocable allocation. The horizon length T can be potentially infinite, which aligns with real-world scenarios where items appear in high volume and frequency, e.g. bandwidth allocation and content recommendation. We assume that agents have linear and additive utilities.

Fair allocation problem in the offline setting has been well-studied. A classical objective to optimize is the *Nash welfare* (NW), defined as the geometric mean of the agents' utilities. Maximizing Nash welfare provides a balance between efficiency and fairness due to the multiplicative nature of the objective. For divisible items, an offline optimal allocation can be computed via solving the Eisenberg-Gale (EG) convex program (Eisenberg and Gale 1959). The solution enjoys both envy-freeness and proportionality, which are important measures of fairness. For indivisible items, finding a Nash welfare maximizing allocation is APX-hard (Moulin 2004; Lee 2017), although constant-factor approximation algorithms are known (Cole and Gkatzelis 2018; Cole et al. 2017; McGlaughlin and Garg 2020).

In the online setting, Gao, Peysakhovich, and Kroer (2021) provides a simple allocation algorithm called PACE (Pace According to Current Estimated utility), which generates asymptotically fair and efficient allocations when items are drawn in an i.i.d. manner. PACE gives each agent a per-round budget of faux currency and simulates a first-price auction in each round. The fair allocation is achieved by having each agent shade their bid with a *pacing multiplier*, which is a projection of their current estimated inverse bang-per-buck to a fixed interval. Liao, Gao, and Kroer (2022) extend these results to non-stationary inputs, where the distribution of items may change over time. They show that in this case, PACE still achieves asymptotic fairness and efficiency guarantees, up to linear error terms from the amount of non-stationarity.

Yet in many real-world scenarios, we cannot expect items to be drawn in a stochastic manner, even if from non-stationary distributions. This motivates the investigation of algorithms with competitive ratio guarantee for adversarial settings. To fit arbitrary inputs, including extreme ones, some algorithms adopt "conservative" designs for fairness such as allocating half of each item purely equally (Banerjee et al. 2022). Although this helps to provide worst-case guarantees, it damages the efficiency in the average case, which may not be acceptable in some practical applications. Moreover, such allocation requires each item to be divisible, or at least for random allocation to be acceptable.

This motivates us to move in another direction: Instead of developing algorithms to fit extreme adversarial inputs, we seek to find worst-case guarantees for existing algorithms that are designed for stochastic inputs, and do not divide any item. In particular, we focus on the performance of the algorithm PACE (Gao, Peysakhovich, and Kroer 2021; Liao, Gao, and Kroer 2022), and explore the question:

How does PACE perform under adversarial input?

Our first contribution is to show that, in the case where all agents have the same weight (or *budget* in market equilibrium terminology), PACE is equivalent to the *first-order integral greedy algorithm*, assuming no projection of the pacing multiplier. Due to this equivalence, we start by investigating the greedy allocation. Our results for first-order

*For the complete paper and technical appendices, see https://arxiv.org/abs/2308.09277.

integral greedy allocation are of independent interest, as it is a natural allocation algorithm.

Although both integral greedy allocation and PACE have infinite envy and $\Omega(T)$ competitive ratio when inputs are completely arbitrary, we notice that such pessimistic results only occur under extreme inputs where the ratio between the largest and smallest non-zero values for an agent differ by an exponential factor. We show that, once we rule out such extreme instances by introducing mild assumptions, both algorithms converge with bounded multiplicative envy competitive ratio *w.r.t.* NW as T increases to infinity. The upper bounds are either constant, or in near-optimal order of n, see Table 1. Combined with existing results under stationary (Gao, Peysakhovich, and Kroer 2021) and non-stationary (Liao, Gao, and Kroer 2022) input models, this establishes a "best-of-many-worlds" guarantee for PACE: it is the first online algorithm that simultaneously guarantees asymptotic fairness and efficiency guarantees under stochastic, stochastic but nonstationary, and adversarial inputs. Plus, as an integral algorithm, PACE is competitive to the optimal continuous allocation. As such, we believe our results show that PACE is a natural and robust algorithm for online fair allocation in real-world settings, since it achieves strong guarantees under many different utility models, and is thus likely to perform well on a variety of real-world inputs.

1.1 Related Work

We review previous works that are most closely related to ours. An extensive review is provided in Appendix A.

The PACE Algorithm. Our work is a direct generalization of the PACE (Pace According to Current Estimated utility) algorithm (Gao, Peysakhovich, and Kroer 2021; Liao, Gao, and Kroer 2022) to adversarial inputs. We will review PACE in detail in Section 2.

Online Fair Division. For maximizing Nash welfare in the online setting, the trivial $\Omega(T)$ competitive ratio result under arbitrary input (Banerjee et al. 2022) motivates the investigation of extra assumptions. Azar, Buchbinder, and Jain (2016) adopted an assumption that we also make: that the minimum *nonzero* valuation of each agent is at least ε times the largest. However, their analysis does not remove the dependence on T in their upper bound of the competitive ratio $O\left(\log \frac{nT}{\varepsilon}\right)$, meaning that the ratio is infinite with an unbounded horizon. Banerjee et al. (2022) assumes extra prior knowledge of the monopolistic utility of each agent, providing an $O(\log n)$ and $O(\log T)$-competitive algorithm; their algorithm involves allocating half of each item uniformly across agents; since we are interested in algorithms with asymptotic convergence guarantees on non-adversarial inputs, such an approach cannot be used. Huang et al. (2022) assumes the input to be λ-balanced or μ-impartial, where λ and μ characterize the desired properties of the input; their competitive ratio upper bounds are logarithmic in the parameter and n. However, their parameters are either scale-variant or implicitly dependent on T. In contrast to the above works, our paper focuses on an asymptotic bound that is free of T as $T \to \infty$. Moreover, while their algorithms can only deal with divisible items, our paper adopts integral allocation, and show that integral decision is sufficient for convergence results given our assumptions.

There is also a line of works on online algorithms with envy guarantees by Bogomolnaia, Moulin, and Sandomirskiy (2022); He et al. (2019); Benade et al. (2018); Zeng and Psomas (2020); Caragiannis et al. (2019), see a detailed review in Appendix A. We note that none of their algorithms achieved convergence in stochastic inputs.

Online Allocation with Resource Constraints. We briefly discuss how this paper differs from existing works in *online resource allocation*, where a sequence of *requests* arrive over time, each consisting of a reward and cost function, and at each time step the algorithm makes a decision to maximize the total rewards subject to long-term cost constraints on each resource. In that setting, strong best-of-many-worlds guarantees are known (Balseiro, Lu, and Mirrokni 2023; Celli, Castiglioni, and Kroer 2022; Castiglioni, Celli, and Kroer 2022). Notably, their objective function is separable across timesteps, e.g., in the form of $\sum_{t=1}^{T} f_t(x)$. This is crucial for regret bounds in these works, as it enables translating dual regret to primal regret by weak duality. However, time-separability no longer holds for Nash welfare; our results cannot be derived with similar techniques to those papers. The types of competitive-ratio guarantees achieved, e.g. by Balseiro, Lu, and Mirrokni (2023), are also impossible in online fair division, where hard input sequences are known (Gao, Peysakhovich, and Kroer 2021; Banerjee et al. 2022).

2 Setup

2.1 Online Fair Allocation

Consider a problem instance with n agents and T items. For $i \in [n]$ and $t \in [T]$, let $v_i^t \geq 0$ be agent i's value for a unit of item t. The input of our problem is a sequence of agent valuations $\boldsymbol{v} = (v_i^t)^{n\times T}$. We assume that for each item j at least one agent values it, i.e. there exists an agent i such that $v_i^t > 0$. Each agent $i \in [n]$ has a non-negative weight B_i, which can also be interpreted as a *budget* of faux currency in a Fisher market (Varian 1974). An allocation $\boldsymbol{x} = (x_i^t)^{n\times T}$ distributes each item to an agent, where x_i^t is the amount of item t that is allocated to agent i gets item t. We assume each item has a unit supply. An allocation is *feasible* if $\sum_{i\in[n]} x_i^t \leq 1$ for each t. An allocation is *integral* if $x_i^t \in \{0, 1\}$. For a feasible, integral allocation x, let $A_i = \{t : x_i^t = 1\}$ be the set of items allocated to agent i.

We assume additive, linear utility for all agents. That is to say, $U_i^t = \sum_{s=1}^{t} x_i^s v_i^s$, where U_i^t is the utility agent i derives from the first t items. For a subset of items $A \subseteq [T]$, agent i's total value on the bundle A is denoted as $U_i(A) = \sum_{s\in A} v_i^s$. Agent i's *monopolistic utility* V_i is defined as his total value for all items $V_i = U_i([T]) = \sum_{t=1}^{T} v_i^t$.

We focus on the online setting where items arrive sequentially, while the set of agents is fixed. An online allocation algorithm makes an irrevocable choice to distribute the item in each round based only on the information of past rounds. Concretely, it maps the history $\mathcal{H}^t = \left\{(v_i^s)_{s=1}^t, (x_i^s)_{s=1}^{t-1}\right\}_{i=1}^n$ to a decision $(x_1^t, \cdots, x_n^t)$ such that $\sum_{i=1}^{n} x_i^t = 1$.

In this paper, we are interested in *envy* and *Nash welfare* (NW) as measures of fairness. The multiplicative envy of

Algorithm	Assumptions	Measure	Upper-Bound ($T \to \infty$)	Theorem
First-order Integral Greedy	Assumption 3.1	multiplicative envy	$1 + 2\log\frac{1}{\varepsilon}$	Theorem 3.4
		competitive ratio *w.r.t.* NW	$\lambda \cdot (n!)^{1/n+\alpha}, \forall \alpha > 0^{*}$	Theorem 3.6
	seed utility δ	utility ratio with seeds	$O(\log T)$	Theorem 3.8
PACE	Assumption 4.1	multiplicative envy	$1 + 2\log\frac{1}{\varepsilon}$	Theorem 4.3
	Assumption 4.2	competitive ratio *w.r.t.* NW	$\left(1 + 2\log\frac{1}{\varepsilon}\right)\frac{1}{c}$	Theorem 4.5

* The lower bound for any online algorithm is at least $(n!)^{1/n}$ when $n \to \infty$.

Table 1: Summary of Results

agent i to agent j is defined as the ratio between the utility that agent i would get from the allocation x_j of agent j to the utility of their own allocation x_i, adjusted by their respective budget. As a criterion for fairness, it measures the extent to which an agent prefers someone else's bundle to his own.

$$\text{Envy}_{ij} = \frac{B_i}{B_j}\frac{\sum_{t=1}^{T} x_j^t v_i^t}{\sum_{t=1}^{T} x_i^t v_i^t}.$$

Notice that when the allocation is integral, the above definition becomes $B_i U_i(A_j)/B_j U_i(A_i)$.

The Nash welfare (NW) of an allocation is defined as the weighted geometric mean of all agents' utilities:

$$\text{NW} = \prod_{i=1}^{n} \left(U_i^T\right)^{B_i/\sum B_j}.$$

Maximizing NW is also equivalent to solving the Eisenberg-Gale convex program of a Fisher market. For a Nash welfare maximizing allocation $\boldsymbol{x}^\star = \{x_i^{\star,t}\}$, let $U_i^\star = \sum_{t=1}^{T} x_i^{\star,t} v_i^t$ be the utility of agent i. (Notice that the optimal allocation might not be unique.) In this paper, we measure the performance in terms of NW maximization based on the *competitive ratio* of our allocation, which is defined as the supremum of the ratio between the online allocation and an optimal offline NW-maximizing allocation, over all possible inputs.

While we allow the input sequence to be adversarial, we restrict our attention to a subset of adversarial input sequences, where we use $\mathcal{V}^T$ to denote the space of valid length T input sequences. Concrete assumptions on $\mathcal{V}^T$ will be specified in Section 3 and Section 4 before analyzing specific algorithms.

We consider the asymptotic worst-case envy and competitive ratio over all possible inputs $\mathcal{V}^T$ when $T \to \infty$:

$$\lim_{T\to\infty} \sup_{v\in\mathcal{V}^T} \max_{i,j} \text{Envy}_{ij}, \; \lim_{T\to\infty} \sup_{v\in\mathcal{V}^T} \prod_{i=1}^{n} \left(\frac{U_i^T}{U_i^\star}\right)^{B_i/\sum B_j}$$

In our analysis, we will show that with proper assumptions, both measures converge with an asymptotic upper bound which is independent of T. Our upper bounds will be either constant or with a near-optimal order dependence on n.

We emphasize that both measures in this paper are defined as ratios, not differences. This is mainly because of the *scale invariant* property of fair division, which is a key feature for desirable allocation algorithms: if an agent's values for all items are multiplied by the same factor, the resulting allocation stays the same. The algorithms that we are interested in, together with NW-maximizing allocations, are all scale invariant. Hence, it is more useful to adopt multiplicative measures, which are also invariant to valuation scaling.

2.2 Algorithms

We introduce the two major algorithms that we study in this paper: the PACE (Pace According to Current Estimated Utility) algorithm (Gao, Peysakhovich, and Kroer 2021; Liao, Gao, and Kroer 2022), and the first-order integral greedy algorithm. Moreover, we will discuss the greedy-based nature of PACE by showing its equivalence to the first-order integral greedy algorithm under certain conditions.

Pseudocode for PACE is shown in Algorithm 2.1. In each round t, the agent utilities are revealed. Each agent then places a bid for the item, which is equal to their value for the item multiplied by the current pacing multiplier β_i^t. The whole item is allocated to the highest bidder, preferring the bidder with the smallest index when a tie occurs. Each agent then observes their realized utility at this round, and updates their current estimated utility. The pacing multiplier is updated to be the weight B_i divided by the estimated utility, then projected to an interval $[a, b]$.

Performance Under Stochastic Input As shown by Gao, Peysakhovich, and Kroer (2021), PACE is an instantiation of stochastic unregularized dual averaging (Xiao 2009) applied to the dual of the underlying allocation program where the supplies are given by the density of each item. With i.i.d. input, PACE converges to the equilibrium of a potentially infinite-dimensional Fisher market (Gao and Kroer 2023), which is closely related to the game-theoretic solution concept of first-price pacing equilibrium (Conitzer et al. 2022). The agent utilities also converge to those associated with the offline NW-maximizing allocation in the mean-square sense.

Theorem 2.1. (Theorem 4. in Gao, Peysakhovich, and Kroer (2021)) *Let $u_i^\star := U_i^\star/T$ be agent i's time-averaged utility under the Nash-welfare-maximizing allocation with supplies given by some underlying distribution. When the values in each rounds are i.i.d. chosen from a distribution, it holds that*

$$\boldsymbol{E}\left[\sum_{t=1}^{T}(\bar{u}_i^t - u_i^\star)^2\right] \le C \cdot \frac{\log T}{T},$$

Algorithm 2.1: PACE

Input: number of agents n, time horizon T, truncation parameters a and b.
Initialization: $U_i^0 = 0, \beta^1 = 1^n$.
for $t = 1, \cdots, T$, **do**
  Agent i bids $\beta_i^t v_i^t$.
  The whole item t is allocated to the highest bidder, with arbitrary tie-breaking:

$$i^t := \min \arg\max_{i\in[n]} \beta_i^t v_i^t, \quad x_i^t = \mathbf{1}(i = i^t).$$

  Agent i updates his current estimated utility

$$\bar{u}_i^t = \frac{1}{t} \cdot x_i^t v_i^t + \frac{t-1}{t} \bar{u}_i^{t-1}.$$

  Agent i updates the pacing multiplier

$$\beta_i^{t+1} = \underset{[a,b]}{\mathrm{Proj}} \left(\frac{B_i}{\bar{u}_i^t} \right),$$

  where $\underset{[a,b]}{\mathrm{Proj}}(\beta) = \max\{a, \min\{b, \beta\}\}$ is the projection operation.
end

where C is a constant independent of T.

Liao, Gao, and Kroer (2022) generalizes Theorem 2.1 to non-stationary inputs, which have a stochastic component yet change over time. Particularly, they consider three input types: independent yet adversarially corrupted input, ergodic input, and periodic input, showing that for all three cases $\boldsymbol{E}\|u^T - u^\star\| \to 0$ is still preserved, up to errors due to non-stationarity.

In the stationary and non-stationary cases, mean-square convergence of time-averaged utility implies an asymptotic competitive ratio of 1 *w.r.t.* NW. For both cases, Gao, Peysakhovich, and Kroer (2021); Liao, Gao, and Kroer (2022) also shows that PACE is asymptotically envy-free (up to a non-stationarity error). In this paper, we provide bounds on PACE's performance with adversarial inputs with assumptions. Combined with previous works, this is the first best-of-many-worlds guarantee for online fair allocation with stationary, non-stationary, and adversarial inputs.

While we attempt to take an algorithm for stochastic inputs and show its performance on adversarial input, the other direction seems to be difficult. It is hard for some algorithms that are designed for adversarial inputs to achieve optimality in stochastic scenarios, due to the conservative routines that they adopt. For instance, Banerjee et al. (2022) divides half of the resources equally, which can be undesirable with stochastic input, see Example 2.2.

Example 2.2. *Consider an online scenario with n agents and n types of items $\{\theta_j\}_{j=1}^n$. Agent i's value for a unit of type-j item $v_i(\theta_j) = 1$ if $i = j$, and $v_i(\theta_j) = 0.01$ otherwise. In each round, the item type is drawn i.i.d. from a uniform distribution. In this scenario,*

- *The equilibrium of the underlying Fisher market assigns all type-i items to agent i. PACE converges to this equilibrium.*
- *The algorithm proposed by Banerjee et al. (2022) allocates at most $\left(\frac{1}{2} + \frac{1}{n}\right)$ fraction of type-i item to agent i, which is clearly not optimal.*

Algorithm 2.2: First-order integral greedy algorithm

Input: number of agents n, time horizon T
Initialization: $U_i^0 = 0$ for all i.
for $t = 1, \cdots, T$, **do**
  Observe agent values for item t, and allocates the whole item to agent i^t:

$$i^t := \min \left(\arg\max_{i\in[n]} \frac{B_i v_i^t}{U_i^{t-1}} \right), \quad x_i^t = \mathbf{1}(i = i^t).$$

  Updates current utility $U_i^t = U_i^{t-1} + x_i^t v_i^t$.
end

Greedy Interpretation of PACE. In the standard configuration of online fair allocation, where agents have equal weight $B_i = 1$, we now show that PACE, when the projection of multipliers is disregarded, can be interpreted as greedily maximizing Nash welfare with integral decisions.

To show this, the following optimization program maximizes NW up to round t greedily, given the history of previous $t-1$ rounds. Its decision is integral.

$$\max_{x_i^t \in \{0,1\}} \quad \sum_{i=1}^n B_i \log \left(\sum_{s=1}^{t-1} x_i^s v_i^s + x_i^t v_i^t \right) \text{ s.t. } \quad \sum_{i=1}^n x_i^t = 1$$

The above program allocates the item to the agent i^t that gives the maximum increment to the objective:

$$i^t \in \arg\max_{i\in[n]} B_i \log \left(1 + \frac{v_i^t}{U_i^{t-1}} \right). \tag{1}$$

With equal weights, i^t is the agent i that maximizes v_i^t / U_i^{t-1}; this coincides with the decision of PACE without projections. This interpretation motivates us to first consider PACE (without projections) as a greedy-fashioned algorithm, and then study PACE (with projections) based on the insights derived from greedy decisions. We focus on *first-order, integral greedy algorithm*, or simply *greedy algorithm* in short, shown in Algorithm 2.2.

In each round, Algorithm 2.2 makes a first-order approximation of the logarithm in (1), and makes a integral decision to greedily maximize this approximation. It is equivalent to (1) only when agent weights are equal. For this consideration, in the following discussions we assume equal weights $B_i = 1$, in accordance with the standard setting in the Nash Welfare maximization literature. We will remark on part of our results that can be generalized to unequal weights. Note that the PACE algorithm itself extends to unequal weights.

PACE can be regarded as greedy algorithm using $\hat{U}_i^t$ as current utility, where $\hat{U}_i^t = \max\{\ell t, \min\{rt, U_i^t\}\}$ is the projected utility and ℓ, r are reinterpreted bounds.

3 Analysis of the Greedy Algorithm

In this section we analyze the performance of the first-order integral greedy algorithm under reasonable assumptions. Missing proofs are deferred to Appendix B.

3.1 Assumptions on the Input

We begin with introducing the assumptions on the input space, as well as the necessity of doing so. We focus on input space $\mathcal{V}_\varepsilon^T$, which is parametrized by $\varepsilon \in (0, 1]$:

Assumption 3.1. *For $T > 0$, $\mathcal{V}_\varepsilon^T$ is the set of inputs which satisfy the following requirements for each $i \in [n]$:*

- *Unbounded monopolistic utility: $V_i = \infty$ $(T \to \infty)$.*
- *Non-extreme values: For each $t \in [T]$, $v_i^t \in \{0\} \cup [\varepsilon, 1]$.*

The first requirement helps to avoid allocating nothing to some agent with integral decision. It is further required that the number of nonzero values of each agent is unbounded, making it meaningful to consider the asymptotic sense in T. The second requirement characterize "non-extreme" nature of agent values, as it is equivalent to assuming a constant bound on the ratio of the minimum nonzero item value and the maximum:

$$\frac{\min_t\{v_i^t : v_i^t > 0\}}{\max_t\{v_i^t\}} \geq \varepsilon,\ \forall i \in [n].$$

The equivalence is due to the scale-invariant property.

In absence of these requirements, the worst-case envy and the competitive ratio is infinite, which makes our anaylsis trivial and uninteresting.

Lemma 3.2. *When v_i^t are arbitrarily chosen from $[0, 1]$, even if the first requirement in Assumption 3.1 is satisfied,*

1. (Banerjee et al. 2022) *Any online allocation algorithm has $\Omega(T)$ competitive ratio.*
2. *The greedy algorithm has $\Omega(T)$ worst-case multiplicative envy.*

Proof Sketch. We construct a hard case for envy. Given horizon T, fix agent 1's value to 1 in all rounds. For agent 2, we set $v_2^t = a^{t-T} \leq 1$ for some $a > 2$. In the greedy allocation, agent 1 receives only one item and has total utility 1. This gives $\mathrm{Envy}_{12} = \Omega(T)$ when $T \to \infty$. □

The above hard example features exponential growth in agent 2's valuation. Hard instances with similar spirit for the continuous problem have previously been given by (Banerjee et al. 2022). The vulnerability of online algorithms with such inputs can be explained by their non-anticipating nature: it cannot see the future. It is difficult for online algorithms to distinguish agents that are hard to satisfy in the future, with those who are easily satisfied.

However, such adversarial instances are arguably not natural. For a real-world market, items are usually similar, e.g., all food or ad slots. It is unlikely for a single agent to have exponentially diverging nonzero values on these items. By requiring a bound on the ratio of minimal and maximal nonzero values, Assumption 3.1 rules out such extreme cases. We will show that once the above assumptions are introduced, both multiplicative envy and competitive ratio *w.r.t.* NW of the greedy algorithm are independent of T asymptotically, i.e., converge to a constant (which depends only on n and ε).

3.2 Envy Analysis for Greedy

For envy analysis, we first observe that envy between any pair of agents can be reduced to 2-agent instances by the inductive structure of the greedy allocation.

Lemma 3.3 (Inductive structure of greedy allocation.)**.** *For any n-agent instance $\boldsymbol{v}$ and agent subset $I \subseteq [n]$, define a new instance $\boldsymbol{v}|_I$ obtained by transforming $\boldsymbol{v}$ as:*

- *Remove all agents that are not in I.*
- *Remove all items that are not in $\bigcup_{i\in I} A_i$.*

Then, the resulting allocation is the same for agents in I, when the Algorithm 2.2 is run on $\boldsymbol{v}$ and $\boldsymbol{v}|_I$.

We show that the multiplicative envy of the greedy allocation is upper bounded by 1 plus a logarithmic term in $1/\varepsilon$, which also generalizes to unequal weights.

Theorem 3.4 (Upper Bound for Multiplicative Envy)**.** *Even with unequal weights, for any $i, j \in [n]$,*

$$\sup_{v\in\mathcal{V}_\varepsilon^T} \mathrm{Envy}_{ij} \leq 1 + 2\log 1/\varepsilon + O\left(T^{-1}\right).$$

Proof Sketch. Due to Lemma 3.3 and symmetry, it suffices to consider Envy_{21} in 2-agent inputs. In this sketch, we assume $B_1 = B_2 = 1$ for simplicity.

We transform any 2-agent input by 1) Setting agent 1's valuation for all items in A_2 to zero; 2) Moving all items in A_2 to the beginning of the input sequence, and all items in A_1 to the end. One can show that allocation under the greedy algorithm is invariant to this transformation. Hence, it suffices to consider only transformed inputs, where agent 2 receives his entire share only in beginning R rounds.

To find the worst-case envy for transformed inputs, a question from an adversarial point of view will be: given $(U_1^R, U_2^R) = (0, U)$, how can we design a value sequence for the coming rounds, such that agent 2's total valuation on items over S rounds is maximized, while ensuring that nothing is allocated to agent 2? This can be characterized by an optimization program:

$$\begin{aligned} \max_{v_1^t, v_2^t \in \{0\}\cup[\varepsilon,1]} \quad & \frac{1}{U}\sum_{t=1}^{\infty} v_2^t \\ \text{s.t.} \quad & v_2^t/v_1^t \leq U/U_1^{t-1}, \quad \forall t \geq 1. \\ & U_1^t \geq \sum_{s=1}^{t} v_1^s, \qquad \forall t \geq 1. \end{aligned} \tag{2}$$

Notice that in (2) we re-index the rounds by starting with index 1 at round $R + 1$. We call (2) a *canonical optimization program* for envy maximization, parametrized by U. Observe that v_2^t/v_1^t is upper bounded by $q(U_1^{t-1})$, defined as

$$q(U_1^{t-1}) = \begin{cases} \min\left\{U/U_1^{t-1}, 1/\varepsilon\right\}, & 0 \leq U_1^{t-1} \leq U/\varepsilon \\ 0, & U_1^{t-1} > U/\varepsilon \end{cases}.$$

We can then give an upper-bound of the objective of the canonical optimization program (2),

$$\frac{1}{U}\sum_{t=1}^{\infty} v_2^t = \frac{1}{U}\sum_{t=1}^{\infty} v_1^t \cdot \frac{v_2^t}{v_1^t} \leq \frac{1}{U}\sum_{t=1}^{\infty} v_1^t \cdot q(U_1^{t-1}). \tag{3}$$

Since increment $v_1^t \leq 1$ is infinitely small when $U \to \infty$, one can show that the right hand side of (3) converges to a definite integral asymptotically as $U \to \infty$:

$$\frac{1}{U}\sum_{t=1}^{\infty} v_1^t \cdot q(U_1^{t-1}) \to \frac{1}{U}\int_0^{U/\varepsilon} q(U_1)\mathrm{d}U_1 = 1 + 2\log\frac{1}{\varepsilon}.$$

The convergence rate on the order of $O(1/T)$ is then achieved by more carefully calculating the above upper bound. □

We complement Theorem 3.4 with a lower bound showing that the bound on multiplicative envy is tight for the first-order integral greedy allocation.

Theorem 3.5 (Lower Bound for Multiplicative Envy of the Greedy Algorithm)**.** *For any $i, j \in [n]$, it holds that*

$$\lim_{T\to\infty} \sup_{v\in\mathcal{V}_\varepsilon^T} \mathrm{Envy}_{ij} \geq 1 + 2\log 1/\varepsilon.$$

3.3 Nash Welfare Analysis for Greedy

We give an upper bound on the asymptotic competitive ratio *w.r.t.* Nash welfare for the greedy algorithm.

Theorem 3.6 (Upper Bound for Competitive Ratio)**.** *For input space $\mathcal{V}_\varepsilon^T$ and any given $\alpha > 0$ there exists a constant $\lambda > 0$ (independent of n and T), such that*

$$\lim_{T\to\infty} \sup_{v\in\mathcal{V}_\varepsilon^T} \left(\frac{\prod_{i=1}^n U_i(A_i)}{\prod_{i=1}^n U_i(A_i^\star)}\right)^{1/n} \leq \lambda\cdot(n!)^{\frac{1+\alpha}{n}}.$$

For $\varepsilon = 1$, the above holds with $\lambda = 1, \alpha = 0$.

Proof Sketch Assume without loss of generality that $U_1(A_1) \leq \cdots \leq U_n(A_n)$. The main idea of the proof is to show that $U_i^\star/U_i$ is bounded by $(n-i+1)\cdot i^\alpha$ asymptotically. Suppose this is not true, we show that $x_i^\star$ will include large proportion of x_j ($j < i$), which will lead to contradiction with the optimality of $\boldsymbol{x}^\star$. □

Although Azar, Buchbinder, and Jain (2016) gives an $O(\log(nT/\varepsilon))$ algorithm, we show that $(n!)^{1/n}$ factor is inevitable if one aims to remove the dependence on T. Hence, first-order integral greedy algorithm is near-optimal in terms of n.

Theorem 3.7 (Lower Bound for Competitive Ratio)**.** *Even in the case $\varepsilon = 1$, for any feasible, deterministic online algorithm we have*

$$\left(\frac{\prod_{i=1}^n U_i(A_i)}{\prod_{i=1}^n U_i^\star}\right)^{1/n} \geq (n!)^{1/n}.$$

Proof Sketch. We construct an adaptive adversary, who attempts to make low-utility agents hard to satisfy in the future. Divide the horizon into n phases, each with length T_i, satisfying $T_i/T_{i-1} \to \infty$. The adversary maintains a set of "active agents", initially containing all agents. In each round, only currently active agents see nonzero values. At the end of each phase, the agent in the active set who has lowest utility is eliminated. This results in a competitive ratio of $(n!)^{1/n}$, see the detailed proof in Appendix B.8. □

3.4 Nash Welfare without Assumption 3.1

As an extension for our analysis on the greedy algorithm, we show how it can be adapted when Assumption 3.1 does not hold. Despite the $\Omega(T)$ lower bound, we show that, when each agent begins with a *seed utility* δ, the competitive ratio of the greedy algorithm is of order $O(\log T)$. The performance is measured *w.r.t.* to the ratio of seeded welfare.

The *seeded* greedy algorithm is identical to Algorithm 2.2, except that all agents are given an initial *seed utility* δ, which is taken into account when deciding the winning agent:

$$i^t := \min\left(\arg\max_{i\in[n]} \frac{v_i^t}{\delta + U_i^{t-1}}\right), \quad x_i^t = \mathbf{1}(i = i^t). \quad (4)$$

The full algorithm is presented in the appendix.

For the seeded algorithm, we study the criterion $R_\delta(\boldsymbol{v})$ which is also defined *w.r.t.* the seeded utility:

$$R_\delta(v) = \sup_{\widetilde{U}} \left\{\frac{1}{n}\sum_{i=1}^n \frac{\widetilde{U}_i + \delta}{U_i + \delta}\right\},$$

where the supremum runs over all feasible hindsight allocations, with resulting utility $\widetilde{U}_1, \cdots, \widetilde{U}_n$. Notice that by the AM-GM inequality, R_δ is also competitive *w.r.t.* the geometric mean. However, due to the presence of the seed utility δ, it is not directly comparable to the criterion $(\prod_{i=1}^n U_i^\star/U_i)^{1/n}$.

Theorem 3.8 (Upper Bound of R_δ for Seeded Greedy)**.** *Run the seeded integral first-order greedy algorithm with seed utility δ. For any input $\boldsymbol{v}$ satisfying $v_i^t \in [0, 1]$,*

$$R_\delta(v) \leq 3 + \frac{4}{\delta} + 2\log\left(1 + \frac{1}{\delta}\right) + 2\log T. \quad (5)$$

The ratio is of order $\log T$, similar to the result given by Azar, Buchbinder, and Jain (2016). Unlike our previous results, here the ratio horizon-dependent since we are considering a broader range of inputs beyond Assumption 3.1. Also, there is a trade-off in choosing seed utility δ. Although a larger δ brings a better bound in (5), it is less able to tell us how the algorithm compares to the offline optimal, since $R_\delta(v) \to 1$ as δ grows large.

4 Analysis of PACE

We continue to analyzing the performance of PACE. Missing examples and proofs can be found in Appendix C.

4.1 Assumptions on the Input

To analyze PACE in the adversarial setting, it is necessary to adopt more assumptions than for the first-order integral greedy algorithm. This is because PACE projects the total utility t to a linear interval $[\ell t, rt]$ in its decision. In the stationary and non-stationary setting of Liao, Gao, and Kroer (2022), the expected utility of each agent grows uniformly with time; in that case the projection helps PACE achieve theoretical guarantees on its performance. However, adversarial input may vary drastically over time, which makes the projection operation more problematic, such as allocating nothing to an agent in the worst case. In their most generic form, our assumptions require that each agent achieves infinite utility as $T \to \infty$ and bounded non-zero valuation ratios;

see Assumption 4.1. Notice that this is necessary if we hope to derive any meaningful convergence guarantees.

Assumption 4.1. *For $T > 0$, $\mathcal{V}_\varepsilon^T(\ell, r)$ is the set of inputs which satisfy for each $i \in [n]$:*

- *$U_i(A_i)$ under PACE is unbounded as $T \to \infty$ when the projection bounds are set to be $[\ell, r]$.*
- *Non-extreme values: For each $t \in [T]$, $v_i^t \in \{0\} \cup [\varepsilon, 1]$.*

Since Assumption 4.1 is potentially hard to verify, we next identify sufficient conditions under which PACE guarantees infinite utility for each agent.

Assumption 4.2. *For integer $T > 0$, and $c \in (0, 1]$, $\mathcal{V}_{\varepsilon,c}^T$ is the set of inputs which satisfy:*

- *For each $i \in [n]$, $V_i \geq cT$, that is, the monopolistic utility of each agent under PACE is $\Theta(T)$.*
- *For each $i \in [n]$ and $t \in [T]$, $v_i^t \in \{0\} \cup [\varepsilon, 1]$.*

Assumption 4.2 strengthens Assumption 3.1 by requiring the monopolistic utility of each agent to be linear in T. Intuitively, this matches the linear projection bounds $[\ell t, rt]$ on utility. We will show in Lemma 4.4 that, with appropriate initialization of the projection bounds, Assumption 4.2 with any $c > 0$ leads to infinite agent utilities, and thus implies Assumption 4.1 for some ℓ and r. Conversely, negative instances can be constructed if $o(T)$ agent monopolistic utility is allowed. Furthermore, initializing projection bounds will require knowledge of c and r. Otherwise, there exists extremely bad input that results in zero utility for some agent.

4.2 Envy Analysis for PACE

Next we show that the PACE algorithm, as long as it is appropriately initialized and Assumption 4.1 holds, achieves $1 + 2\log 1/\varepsilon$ multiplicative envy bound asymptotically. Our proof is performed by reducing the adversarial envy-maximizing problem in PACE to the canonical program (2) by showing that PACE with our assumptions lead to a weakly harder problem than (2) for the adversary.

Theorem 4.3 (Upper Bound of Multiplicative Envy for PACE)**.** *Under Assumption 4.1 with $r = 1$ and $\ell < \varepsilon^2 \cdot \{\varepsilon + 1 + (n-1)(1 + \log(1/\varepsilon))\}^{-1}$, PACE achieves*

$$\lim_{T\to\infty} \sup_{\boldsymbol{v} \in \mathcal{V}_\varepsilon^T(\ell,r)} \text{Envy}_{ij} \leq 1 + 2\log 1/\varepsilon, \ \ \forall i, j \in [n].$$

Proof Sketch. Similar to Theorem 3.4, we construct a transformation and impose it on any given input. However, the reduction to 2-agent cases no longer holds. It is also not obvious that the transformation preserves PACE's allocation. Considering these challenges, we elaborate the proof into three major steps:

- First, we show that for some t^*, the PACE allocation is preserved for the first t^* items in the transformed sequence.
- Second, we show that for the first t^* items in the transformed sequence, the problem of maximizing envy under PACE can be reduced into a canonical program, i.e., the constraints are weakly harder than (2). This gives upper bounds on multiplicative envy, as well as the number of items with nonzero agent valuations in the first t^* rounds.
- Finally, with the bound on item numbers, we show that the transformed sequence indeed terminates within t^* rounds.

□

4.3 Nash Welfare Analysis for PACE

Next we focus on establishing a worst-case guarantee for the asymptotic competitive ratio of Nash Welfare for PACE under Assumption 4.2. To begin with, we show that Assumption 4.2 implies infinite agent utilities with appropriate initialization.

Lemma 4.4. *For any input $\boldsymbol{v} \in \mathcal{V}_{\varepsilon,c}^T$, if $r = 1, \ell < \frac{c\varepsilon^2}{1+(n-1)(1+\log 1/\varepsilon)}$ then there exists a constant $d > 0$ which depends on ℓ, such that PACE satisfies $\lim_{T\to\infty} U_i(A_i) \geq dT$ for each i. Furthermore, $d/c = \Omega(n^{-1})$ as $n \to \infty$.*

The proof of Lemma 4.4 is also by a reduction from the canonical problem (2). We remark that Lemma 4.4 yields more than infinite agent utilities: it also tells us that the utilities are linear in T. Moreover, from $d/c = \Omega\left(n^{-1}\right)$ we know that PACE computes an asymptotic approximate proportional allocation, which helps to derive a bounded competitive ratio *w.r.t.* Nash welfare. The bound can be furthermore refined using the envy results.

Theorem 4.5 (Upper Bound of Competitive Ratio for PACE)**.** *For any input $\boldsymbol{v} \in \mathcal{V}_{\varepsilon,c}^T$, if $r = 1, \ell < \frac{c\varepsilon^2}{1+\varepsilon+(n-1)(1+\log 1/\varepsilon)}$, PACE achieves*

$$\lim_{T\to\infty} \sup_{v \in \mathcal{V}_\varepsilon^T} \left(\frac{\prod_{i=1}^n U_i^\star}{\prod_{i=1}^n U_i(A_i)} \right)^{1/n} \leq \left(1 + 2\log\frac{1}{\varepsilon}\right) \cdot \frac{1}{c}.$$

Theorem 4.5 gives an upper bound depends only on parameters ε and c, which are both independent of T. We remark that the constant $1/c$ might not be tight.

We also remark that in both Theorem 4.3 and Theorem 4.5 ℓ is at most the order of $1/n$, which is aligned with the stationary and non-stationary setting (projecting utilities to an $\Omega(1/n)$ bound is unreasonable since there are n agents). However with adversarial input ℓ decreases as $\varepsilon \to 0$, which means that PACE requires a wider projection interval when its input becomes potentially more extreme.

5 Conclusion

We proved horizon-independent bounds for envy and Nash welfare for both the first-order integral greedy algorithm and PACE under adversarial inputs with mild assumptions. Our results complete the first best-of-many-worlds result for online fair allocation, since PACE thus achieves guarantees under stochastic (Gao, Peysakhovich, and Kroer 2021), stochastic but nonstationary (Liao, Gao, and Kroer 2022), and adversarial inputs. Moreover, our results on greedy algorithm are of independent interest, as they characterize assumptions needed to achieve guarantees for that algorithm.

It remains open whether the constant in Theorem 4.5 can be improved. A more general open question is to explore more best-of-many-worlds online fair allocation algorithms, with potentially different performance measures and assumptions.

Acknowledgments

This research was supported by the Office of Naval Research awards N00014-22-1-2530 and N00014-23-1-2374, and the National Science Foundation awards IIS-2147361 and IIS-2238960.

References

Anari, N.; Oveis Gharan, S.; Saberi, A.; and Singh, M. 2017. Nash Social Welfare, Matrix Permanent, and Stable Polynomials. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Azar, Y.; Buchbinder, N.; and Jain, K. 2016. How to Allocate Goods in an Online Market? *Algorithmica*, 74(2): 589–601.

Bach, F.; and Levy, K. Y. 2019. A universal algorithm for variational inequalities adaptive to smoothness and noise. In *Conference on learning theory*, 164–194. PMLR.

Balseiro, S. R.; Lu, H.; and Mirrokni, V. 2023. The best of many worlds: Dual mirror descent for online allocation problems. *Operations Research*, 71(1): 101–119.

Banerjee, S.; Gkatzelis, V.; Gorokh, A.; and Jin, B. 2022. Online Nash social welfare maximization with predictions. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1–19. SIAM.

Barman, S.; Bhaskar, U.; Krishna, A.; and Sundaram, R. G. 2020. Tight approximation algorithms for p-mean welfare under subadditive valuations. *arXiv preprint arXiv:2005.07370*.

Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018. Finding fair and efficient allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 557–574.

Benade, G.; Kazachkov, A. M.; Procaccia, A. D.; and Psomas, C.-A. 2018. How to make envy vanish over time. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 593–610.

Bogomolnaia, A.; Moulin, H.; and Sandomirskiy, F. 2022. On the fair division of a random object. *Management Science*, 68(2): 1174–1194.

Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3): 1–32.

Castiglioni, M.; Celli, A.; and Kroer, C. 2022. Online Learning under Budget and ROI Constraints and Applications to Bidding in Non-Truthful Auctions. *arXiv preprint arXiv:2302.01203*.

Celli, A.; Castiglioni, M.; and Kroer, C. 2022. Best of Many Worlds Guarantees for Online Learning with Knapsacks. *arXiv preprint arXiv:2202.13710*.

Chaudhury, B. R.; Garg, J.; and Mehta, R. 2021. Fair and Efficient Allocations under Subadditive Valuations. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Cole, R.; Devanur, N.; Gkatzelis, V.; Jain, K.; Mai, T.; Vazirani, V. V.; and Yazdanbod, S. 2017. Convex program duality, Fisher markets, and Nash social welfare. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 459–460.

Cole, R.; and Gkatzelis, V. 2018. Approximating the Nash social welfare with indivisible items. *SIAM Journal on Computing*, 47(3): 1211–1236.

Conitzer, V.; Kroer, C.; Panigrahi, D.; Schrijvers, O.; Stier-Moses, N. E.; Sodomka, E.; and Wilkens, C. A. 2022. Pacing equilibrium in first price auction markets. *Management Science*, 68(12): 8515–8535.

Eisenberg, E.; and Gale, D. 1959. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1): 165–168.

Gao, Y.; and Kroer, C. 2023. Infinite-dimensional Fisher markets and tractable fair division. *Operations Research*, 71(2): 688–707.

Gao, Y.; Peysakhovich, A.; and Kroer, C. 2021. Online market equilibrium with application to fair division. *Advances in Neural Information Processing Systems*, 34: 27305–27318.

Garg, J.; Hoefer, M.; and Mehlhorn, K. 2018. Approximating the Nash social welfare with budget-additive valuations. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2326–2340. SIAM.

Garg, J.; Kulkarni, P.; and Kulkarni, R. 2020. Approximating Nash social welfare under submodular valuations through (un) matchings. In *Proceedings of the fourteenth annual ACM-SIAM symposium on discrete algorithms*, 2673–2687. SIAM.

He, J.; Procaccia, A. D.; Psomas, C.; and Zeng, D. 2019. Achieving a fairer future by changing the past. *IJCAI'19*.

Huang, Z.; Li, M.; Shu, X.; and Wei, T. 2022. Online Nash Welfare Maximization Without Predictions. *arXiv preprint arXiv:2211.03077*.

Kaneko, M.; and Nakamura, K. 1979. The Nash social welfare function. *Econometrica: Journal of the Econometric Society*, 423–435.

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

Li, W.; and Vondrák, J. 2022. A constant-factor approximation algorithm for Nash social welfare with submodular valuations. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, 25–36. IEEE.

Liao, L.; Gao, Y.; and Kroer, C. 2022. Nonstationary dual averaging and online fair allocation. *Advances in Neural Information Processing Systems*, 35: 37159–37172.

McGlaughlin, P.; and Garg, J. 2020. Improving Nash social welfare approximations. *Journal of Artificial Intelligence Research*, 68: 225–245.

Moulin, H. 2004. *Fair division and collective welfare*. MIT press.

Nash Jr, J. F. 1950. The bargaining problem. *Econometrica: Journal of the econometric society*, 155–162.

Varian, H. R. 1974. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1): 63–91.

Xiao, L. 2009. Dual averaging method for regularized stochastic learning and online optimization. *Advances in Neural Information Processing Systems*, 22.

Zeng, D.; and Psomas, A. 2020. Fairness-efficiency tradeoffs in dynamic fair division. In *Proceedings of the 21st ACM Conference on Economics and Computation*, 911–912.