Nonstationary Dual Averaging and Online Fair Allocation

Luofeng Liao, Yuan Gao, Christian Kroer
IEOR, Columbia University
{ll3530,yg254,ck294}@columbia.edu

Abstract

We consider the problem of fairly allocating sequentially arriving items to a set of individuals. For this problem, the recently-introduced PACE algorithm leverages the dual averaging algorithm to approximate competitive equilibria and thus generate online fair allocations. PACE is simple, distributed, and parameter-free, making it appealing for practical use in large-scale systems. However, current performance guarantees for PACE require i.i.d. item arrivals. Since real-world data is rarely i.i.d., or even stationary, we study the performance of PACE on nonstationary data. We start by developing new convergence results for the general dual averaging algorithm under three nonstationary input models: adversarially-corrupted stochastic input, ergodic input, and block-independent (including periodic) input. Our results show convergence of dual averaging up to errors caused by nonstationarity of the data, and recover the classical bounds when the input data is i.i.d. Using these results, we show that the PACE algorithm for online fair allocation simultaneously achieves "best of many worlds" guarantees against any of these nonstationary input models as well as against i.i.d. input. Finally, numerical experiments show strong empirical performance of PACE against nonstationary inputs.

1 Introduction

In fair division, the goal is to allocate a set of items, typically assumed divisible, among a set of agents with heterogeneous preferences, while guaranteeing fairness and efficiency properties. In this paper we are interested in how to fairly and efficiency allocate items that arrive *online*: at every time step one item arrives, and we must irrevocably assign it to some agent. Recently, there has been a growing literature on such online fair allocation problems (Azar et al., 2016; Balseiro et al., 2020; Gao et al., 2021; Bateni et al., 2021; Sinclair et al., 2021; Banerjee et al., 2022). Real-world systems that can be captured by such settings include Internet advertising systems, job recommender systems, cloud computing platforms, and many more. One of the key challenges in such problems is to balance the (often conflicting) goals of overall efficient resource utilization with fairness guarantees for the individual agents.

For this setting, Gao et al. (2021) shows that a simple mechanism called PACE (Pace According to Current Estimated utility) generates asymptotically fair and efficient allocations when the item arrivals are drawn in an i.i.d. manner. PACE gives each agent a per-time-step budget of faux currency, and the fair allocation is achieved by having agents participate in first-price auctions for each item, using the faux money. By guaranteeing that each agent asymptotically spends their budget at the correct rate, the resulting allocations and prices converge to what is known as a *competitive equilibrium from equal incomes* (CEEI), which guarantees both fairness and efficiency. In PACE, each agent maintains a *pacing multiplier* to control their spending over time, and the pacing multipliers are updated based on buyers' budgets and cumulative utilities. This is similar to how budget-management

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

systems work in Internet ad auctions. PACE is highly decentralized due to its auction-based allocation, it does not require dividing the item, and it is also completely parameter free. This makes it suitable for large-scale practical implementation.

Yet in many large-scale settings, such as the context of fair recommender systems (Kroer et al., 2021; Kroer and Stier-Moses, 2022) or Internet advertising, we would not expect items to be drawn i.i.d. from a single distribution. One alternative is to assume that data arrives adversarially. However, this leads to very pessimistic negative results and is not an accurate representation of the data one would expect to see in practice. Instead, one would expect the data to have a strong stochastic component, but with changes over time, e.g., due to flow of traffic, breaking news events, or system updates (Esfandiari et al., 2018; Balseiro et al., 2020).

Motivated by the above considerations, we study online fair allocation when the data exhibits nonstationary behavior. In particular, we focus on the performance of the PACE algorithm of Gao et al. (2021). We ask

How does PACE behave when nonstationarity is present in the stream of items?

We show that, under several data-input models, the fairness and efficiency guarantees of the PACE algorithm are still preserved, up to errors due to the nonstationarity of the data input. In this sense, we significantly extend the main results in Gao et al. (2021). To show these results, we first consider the more general setting of nonstationary stochastic optimization and develop new performance guarantees for dual averaging in this setting. Given the ubiquitous use of dual averaging in online and stochastic optimization, our results are of broader interest beyond (fair) resource allocation.

1.1 Summary of Contributions

Novel convergence results for dual averaging under three nonstationary settings.

We analyze the dual averaging (DA) algorithm for nonstationary stochastic optimization under different data input models, namely, (1) mildly corrupted, (2) ergodic and (3) periodic input data. Specifically, we consider the composite dual averaging algorithm, where the composite term is strongly convex. We show that, in all cases, the iterates generated by dual averaging (DA) converge to the optimal solution in mean square, where the bound on the mean-square error decomposes into two terms: i) the typical $O(\log t/t)$ guarantee known from the i.i.d. case, and ii) a term that depends on the amount of nonstationarity in the data input model. Our results recover the classical bounds under i.i.d. data input as a special case.

Theoretical fairness and efficiency guarantees of PACE for nonstationary item arrivals.

We consider the online fair allocation problem where item arrivals follow any of the three data input models that we consider for DA; these settings generalize the i.i.d. setting in Gao et al. (2021). Utilizing our convergence results for DA under nonstationary data input models, we show that, for item arrivals following these models, PACE ensures convergence of the pacing multipliers, again with a decomposition into a $O(\log t/t)$ term as well as a term depending on the nonstationarity. We then show that the agents' realized utilities, envy, regrets, and expenditures all obtain convergence bounds based on the convergence of pacing multipliers. Our results show that PACE as an online fair resource allocation algorithm is robust against distributional uncertainty of the input and automatically adapts to many different data input models without any parameter tuning. In Appendix F we provide numerical experiments which corroborate the above theory and demonstrate the practical efficiency of PACE under different data input models.

An extensive review of related work is provided in Appendix A.

Notation. We use 1_t to denote the vector of ones of length t and e_j to denote the vector with one in the j-th entry and zeros in the others. We use $\Delta(\Theta)$ to denote the space of probability measures on a measurable space Θ, and Δ_n to denote the simplex in $\mathbb{R}^n$. To measure the nonstationarity in the input data, we will use the total variation distance. Given two probability measures P and Q, it is defined as $\|P - Q\|_{\mathrm{TV}} := (1/2) \int |\frac{\mathrm{d}P}{\mathrm{d}\mu} - \frac{\mathrm{d}Q}{\mathrm{d}\mu}| \, \mathrm{d}\mu$, where μ is a supporting measure.

2 Preliminaries on Online Fair Allocation

An online fair allocation instance with infinitely divisible items with n agents and a finite horizon t consists of a tuple $\mathsf{A} = (n, t, \Theta, Q, v)$, where Θ is the (possibly uncountable) measurable space of all possible items, with an associated σ-algebra $\mathcal{M}$ and a probability measure μ, the distribution $Q \in \Delta(\Theta^t)$ is the distribution over possible sequences of items $\gamma = (\theta_1, \ldots, \theta_t) \in \Theta^t$, each of unit supply, and the set $v = (v_1, \ldots, v_n) \in L^1_+(\Theta)^n$ is the set of valuation functions of the n agents. Here $L^1_+(\Theta)$ is the space of positive integrable functions on Θ. Agent i sees a utility of $v_i(\theta)$ in item $\theta \in \Theta$. Abusing notation we let $v_i(\gamma) = \big(v_i(\theta^1), \ldots, v_i(\theta^t)\big)$ denote the valuation for agent i of items in the sequence γ. Let Q^τ be the marginal distribution of the item θ^τ at time τ and $\bar{Q} = (1/t)\sum_{\tau=1}^t Q^\tau$. We assume $\int_\Theta v_i \mathrm{d}\bar{Q} = 1$ for all $i \in [n]$. We further assume $|v|_\infty := \max_i \|v_i\|_\infty < \infty$. We stress that the PACE algorithm that we study is not going to require access to either the valuation functions v or the set of possible items Θ; these are only required in order to discuss the resulting bounds.

Given an instance A, the decision maker allocates the stream of items γ one at a time, in an irrevocable manner. At time τ when item θ_τ is revealed, the decision maker must choose an allocation $x^\tau = (x_1^\tau, \ldots, x_n^\tau) \in \Delta_n$ based on information available at that time, and allocate accordingly. Here the i-th entry of x^τ is the fraction of item θ^τ allocated to agent i. On receiving her fraction, agent i realizes a utility of $u_i^\tau := v_i(\theta^\tau)x_i^\tau$. We let $x = (x^1, \ldots, x^t)$ denote the sequence of allocations made over time. For agent i, let $x_i = (x_i^1, \ldots, x_i^t) \in \mathbb{R}^t$ denote the fraction of items given to agent i across time. With this notation, the total utility of agent i is $\langle x_i, v_i(\gamma)\rangle$. The goal of the decision maker is to choose, in an online fashion, an allocation x such that it achieves some form of both efficiency and fairness guarantees.

2.1 Benchmark: The Hindsight Allocation

As a benchmark, we will consider the hindsight-optimal allocation. Suppose all items are presented to the decision maker as opposed to arriving one by one. In that case, a fair and efficient allocation can be found by allocating using the *Eisenberg-Gale* (EG) convex program Eisenberg and Gale (1959). EG picks the allocation that maximizes the sum of weighted logarithmic utilities (which is equivalent to maximizing the weighted geometric mean of utilities):

$$\max_{x\geq 0, u\geq 0} \left\{ \sum_{i=1}^n B_i \log(U_i) \;\middle|\; U_i \leq \big\langle v_i(\gamma), x_i \big\rangle \;\; \forall i \in [n] \;,\; \sum_{i=1}^n x_i^\tau \leq 1 \;\; \forall \tau \in [t] \right\}. \tag{1}$$

The weights B_i represent the priority given to each agent, and they they can be interpreted as budgets in a market-based interpretation of the EG allocation.[1]. We will focus on the case where $B_i = t/n$ for all i, but all our results extend directly to the case of unequal weights, which can be useful in settings such as when buyers have quasilinear utilities Gao et al. (2021); Conitzer et al. (2019) or when it is desirable to give a larger allocation to certain agents.

The PACE algorithm asymptotically converges to the optimal dual solution, which is

$$\beta^\gamma := \operatorname*{arg\,min}_{\beta\geq 0} \left\{ \frac{1}{t}\sum_{\tau=1}^t \max_{i\in[n]} \beta_i v_i(\theta^\tau) - \frac{1}{n}\sum_{i=1}^n \log \beta_i \right\}. \tag{2}$$

We will also be interested in the *underlying problem* implied by the average item supplies $s = \mathrm{d}\bar{Q}/\,\mathrm{d}\mu$. Letting $\big\langle v_i, x_i \big\rangle := \int_\Theta v_i x_i \,\mathrm{d}\mu$, this leads to the infinite-dimensional analogue of (1):

$$\max_{x\in L^\infty_+(\Theta), u\geq 0} \left\{ \frac{1}{n}\sum_{i=1}^n \log(u_i) \;\middle|\; u_i \leq \big\langle v_i, x_i \big\rangle \;\; \forall i \in [n], \;\; \sum_{i=1}^n x_i \leq s \right\}, \tag{3}$$

We let u^* denote the optimal utilities in Eq. (3). The infinite-dimensional analogue of the dual (2) is the following. For any $\delta_0 > 0$, The infinite-dimensional analogue of (2) is the following. For any $\delta_0 > 0$,

$$\beta^* := \operatorname*{arg\,min}_{\frac{1}{n(1+\delta_0)}\leq\beta\leq 1+\delta_0} \left\{ \int_\Theta \Big(\max_{i\in[n]} \beta_i v_i(\theta)\Big)\bar{Q}(\mathrm{d}\theta) - \frac{1}{n}\sum_{i=1}^n \log \beta_i \right\}. \tag{4}$$

[1]The hindsight allocation Eq. (1) can be interpreted as a competitive equilibrium from equal incomes (CEEI) in the corresponding Fisher market; see Appendix B for more details on this interpretation.

A rigorous mathematical treatment of the infinite-dimensional program can be found in Gao and Kroer (2021) and (Gao et al., 2021, Section 2). Note the additional constraint in Eq. (4) on β does not affect the optimal solution since $1/n \leq \beta_i^* \leq 1$; see Lemma 1 in Gao and Kroer (2021).

It is well-known that the hindsight allocation generated by the EG program enjoys the following efficiency and fairness properties:

1. Pareto optimality: we cannot strictly increase any agent's utility without decreasing some other agents' utility.
2. Envy-freeness: each agent prefers their own allocation to that of any other agent: $\langle v_i(\gamma), x_i^* \rangle \geq \langle v_i(\gamma), x_k^* \rangle$ for all $k \neq i$.
3. Proportionality: every agent achieves at least as much utility as under the uniform allocation, i.e. $\langle v_i(\gamma), x_i^* \rangle \geq \langle v_i(\gamma), (1/n)1_t \rangle$.

Therefore the hindsight EG allocation is the gold standard that we assume the decision maker would use if she had known the sequence of items γ in advance. However, in the online setting the decision maker does not know this sequence, and must therefore instead attempt to approximate an equally good allocation in online fashion.

For an item sequence γ, we let x^γ denote the optimal hindsight allocation, which is an optimal solution to Eq. (1), and we denote the resulting total and average utility as

$$U_i^\gamma := \left\langle v_i(\gamma), x_i^\gamma \right\rangle = \sum_{\tau=1}^{t} x_i^{\gamma,\tau} v_i(\theta^\tau), \quad u_i^\gamma := (1/t) \cdot U_i^\gamma \ . \tag{5}$$

2.2 Performance Metrics

We measure the performance of an online allocation rule x on the instance γ via the following two quantities. The *regret* of agent i is the difference between the total hindsight equilibrium utility U_i^γ and their realized utilities u_i^τ under x

$$\mathrm{Reg}_{i,t}(\gamma) := U_i^\gamma - \sum_{\tau=1}^{t} u_i^\tau \ . \tag{6}$$

The *envy* of agent i is the maximal extent to which they prefer the allocation of any other agent:

$$\mathrm{Envy}_{i,t}(\gamma) := \max_{k \in [n]} \left\{ \langle v_i(\gamma), x_k \rangle - \langle v_i(\gamma), x_i \rangle \right\} \ . \tag{7}$$

We seek to understand the worst-case behavior of an algorithm when facing a certain class of input distributions. For a given input distribution class $\mathsf{C} \subset \Delta(\Theta^t)$, we will develop bounds on the worst-case regret and envy under any distribution in C:

$$\sup_{Q \in \mathsf{C}} \mathbb{E}_{\gamma \sim Q} \left[\mathrm{Reg}_{i,t}(\gamma) \right] \ , \quad \sup_{Q \in \mathsf{C}} \mathbb{E}_{\gamma \sim Q} \left[\mathrm{Envy}_{i,t}(\gamma) \right] \ .$$

2.3 The PACE Algorithm

In this section, we review the PACE (Pace According to Current Estimated Utility) dynamics (Gao et al., 2021). In PACE, each item is allocated via first-price auction, and each agent constructs bids by scaling their value by a *pacing multiplier*. The pacing multipliers are maintained using simple, distributed updates that can be handled either by the agents or by the platform.

Algorithmic details are displayed in Algorithm 1. Here $\Pi_{[\ell,h]}[x] = \max\{\ell, \min\{x, h\}\}$. At every time step τ an item θ^τ is revealed. At that point every agent comes up with a *bid* for that item, which is equal to their value for the item multiplied by their current pacing multiplier β_i^τ. Then, the agents submit these bids to a first-price auction, and the item is allocated to the highest bidder. For concreteness, we choose the bidder with the smallest index if a tie occurs, but any rule works. Each agent then observes their realized utility, updates their average utility received so far, and updates their pacing multiplier accordingly. As pointed out in Gao et al. (2021), PACE is an instantiation of dual averaging Xiao (2010) applied to the dual of the hindsight allocation program in (4).

PACE has many attributes desirable in real-world applications.

Algorithm 1: PACE(n, t, δ_0)

Input: number of agents n, horizon t, algorithm parameter $\delta_0 > 0$.

Initialize: Set $\beta^1 = (1 + \delta_0) \cdot 1_n$.

Environment draws the item sequence $\gamma = \{\theta^1, \ldots, \theta^t\}$ from the distribution Q.

for $\tau = 1, \ldots, t$ *when item* θ^τ *is revealed* **do**

  Agent i bids $\beta_i^\tau v_i(\theta^\tau)$, the whole item θ^τ is allocated to the highest bidder i^τ (with arbitrary tie breaking) $i^\tau := \min\{\arg\max_{i \in [n]} \beta_i^\tau v_i(\theta^\tau)\}$.

  Agent i updates current average utility $u_i^\tau = v_i(\theta^\tau)\mathbb{1}\{i = i^\tau\}$, $\quad \bar{u}_i^\tau = \frac{1}{\tau}\sum_{s=1}^{\tau} u_i^s$.

  Agent i updates the pacing multiplier $\beta_i^{\tau+1} = \Pi_{[\ell,h]}\left[1/(n\bar{u}_i^\tau)\right]$, where the interval $[\ell, h] = \left[\frac{1}{(1+\delta_0)n}, 1 + \delta_0\right]$.

Highlight 1. Decentralization. The PACE dynamics can be run in either centralized (by having the mechanism designer emulate the pacing process for each agent) or decentralized fashion (since the auction-based allocation is the only centralized step at each iteration), and are therefore suitable for Internet-scale online fair division and online Fisher market applications.

Highlight 2. Pure Allocation. PACE allows each item to be fully allocated to a single agent, even though the hindsight performance metric is allowed to utilize fractional allocations. While fractional allocations can be interpreted as randomized allocations in many large-scale settings, this may not always be desirable, for example when allocating food to food banks.

Highlight 3. Tuning-free. An important fact about the PACE dynamics is that each agent has no stepsize parameter whatsoever, which means that no stepsize tuning is required. Moreover, PACE is robust against *the types* of item arrivals since the algorithm needs neither knowledge of the item distribution P nor the input type C.

In addition to the regret and the envy performance metrics, we will also derive results for the following two quantities that characterize the long-run behavior of PACE. Let $\bar{u}^t = (1/t) \cdot \sum_{\tau=1}^{t} u^\tau$ be the vector of average realized utilities for all agents. We will show that the agents' utilities converge to those associated to the *underlying offline fair allocation problem*, u^*, defined in Eq. (3), in a mean-square sense, i.e., $\mathbb{E}\left[\|\bar{u}^t - u^*\|^2\right] \to 0$, as long as the error due to nonstationarity grows sublinearly in the number of time periods. Secondly, define the *expenditure* of agent i at time τ by $b_i^\tau := \beta_i^\tau v_i(\theta^\tau)\, \mathbb{1}\{i = i^\tau\}$. We will show $(1/t) \cdot \sum_{\tau=1}^{t} b_i^\tau \to 1/n$ in mean square as well, as long as the error due to nonstationarity grows sublinearly in the number of time periods.

3 Main Results

This section introduces the main results of this paper: the behavior of PACE under three different types of nonstationary input models. All prior results on fair online allocation have been either for worst-case inputs (with much more conservative guarantees and not for the PACE algorithm) (Azar et al., 2016; Banerjee et al., 2022) or for i.i.d. input data Gao et al. (2021).

We first introduce some notation that will be useful for describing these input models. For $s > \tau \geq 1$ let $Q^s(\theta^{1:\tau})$ denote the conditional distribution of θ^s given $\{\theta^1, \ldots, \theta^\tau\}$. For a subset I of $[t]$ let Q^I denote the joint distribution of the variables $\{\theta^\tau\}_{\tau \in I}$. Let $\bar{Q} = (1/t) \cdot \sum_{\tau=1}^{t} Q^\tau$ be the uniform mixture of $\{Q^\tau\}_\tau$. We study three types of input: independent input with adversarial corruption, ergodic and Markov input, and periodic input. For each input setting, we describe our main theorem for the performance guarantees of PACE here. The proofs are given in Appendix C, because these results rely on developing a theory of nonstationary performance of DA, which is done in Section 4.

3.1 Independent Input with Adversarial Corruption

Adversarial perturbation of a fixed item distribution models scenarios where the items generally behave in a predictable manner, but for some time steps the input behaves erratically. Typically this is assumed to happen only for a small number of time steps. Such perturbation could be malicious,

for example when item arrivals are manipulated in favor of certain agents; or non-malicious, such as unpredictable surges of certain keywords on search engines (Esfandiari et al., 2018).

We study a type of adversarial perturbation where the item distribution at each time step might be corrupted by an arbitrary amount, but distributions at different time steps are independent of each other. We assume the average corruption is bounded by δ, as measured in TV distance. The set of distributions over sequences that we consider is then:

$$\mathsf{C}^{\mathrm{ID}}(\delta) := \left\{ Q \in \Delta(\Theta)^t : \frac{1}{t}\sum_{\tau=1}^{t} \|Q^\tau - \bar{Q}\|_{\mathrm{TV}} \le \delta \right\} . \tag{8}$$

We use $\tilde{O}$ to hide numeric constants and polynomials of n, $|v|_\infty$, and $\log t$. Our main fair online allocation result for the adversarial corruption case is:

Theorem 1 (Independent Case)**.** *For the adversarially corrupted and independent case, Algorithm 1 guarantees that for an instance* $\mathsf{A} = (n, t, \Theta, Q, v)$*, we have*

$$\sup_{Q \in \mathsf{C}^{\mathrm{ID}}(\delta)} \mathbb{E}_{\gamma\sim Q}\big[\mathrm{Reg}_{i,t}(\gamma)\big] \ , \ \sup_{Q \in \mathsf{C}^{\mathrm{ID}}(\delta)} \mathbb{E}_{\gamma\sim Q}\big[\mathrm{Envy}_{i,t}(\gamma)\big] = \tilde{O}\big(\sqrt{t} + \sqrt{\delta}\cdot t\big) \tag{9}$$

and

$$\sup_{Q \in \mathsf{C}^{\mathrm{ID}}(\delta)} \mathbb{E}_{\gamma\sim Q}\left[\|\bar{b}^t - (1/n)1_n\|^2\right] \ , \sup_{Q \in \mathsf{C}^{\mathrm{ID}}(\delta)} \mathbb{E}_{\gamma\sim Q}\left[\|\bar{u}^t - u^*\|^2\right] \ , \sup_{Q \in \mathsf{C}^{\mathrm{ID}}(\delta)} \mathbb{E}_{\gamma\sim Q}\left[\|\bar{u}^t - u^\gamma\|^2\right] = \tilde{O}(\delta + 1/t) \ .$$

The result shows that the regret and envy performance metrics degrade linearly in the average corruption δ. In the i.i.d. case where $\delta = 0$, we recover the $\sqrt{t}$ regret rate and the $1/t$ rate of convergence for utilities and expenditures in terms of the mean-square error from Gao et al. (2021). If out of the t distributions of items in each time step only $O(\sqrt{t})$ are corrupted, each by a constant amount, then the $\sqrt{t}$ regret and envy bounds, as well as $1/t$ convergence rates, are also preserved.

3.2 Ergodic Input and Markov Processes

To handle correlation across time, we next study ergodic inputs. For these inputs, strong correlation might be present for items sampled at nearby time steps, but the correlation between items decays as they are separated in time. For any integer ι such that $1 \le \iota \le t-1$, we measure the ι-step deviation from some distribution $\Pi \in \Delta(\Theta)$ by the quantity

$$\delta(\iota) := \sup_{\gamma} \sup_{\tau=1,\ldots,t-\iota} \|Q^{\tau+\iota}(\theta^{1:\tau}) - \Pi\|_{\mathrm{TV}} \ .$$

Intuitively, this definition tells us that, no matter where and when we start the item arrival process, it takes only ι steps to get $\delta(\iota)$-close to the distribution Π. We will consider the set of ergodic input distributions whose ι-step deviation is bounded by δ:

$$\mathsf{C}^{\mathrm{E}}(\delta, \iota) := \left\{ Q \in \Delta(\Theta^t) : \sup_{\gamma} \sup_{\tau=1,\ldots,t-\iota} \|Q^{\tau+\iota}(\theta^{1:\tau}) - \Pi\|_{\mathrm{TV}} \le \delta, \text{ for some } \Pi \in \Delta(\Theta) \right\} . \tag{10}$$

Theorem 2 (Ergodic Case)**.** *For the ergodic case, Algorithm 1 guarantees that for an instance* $\mathsf{A} = (n, t, \Theta, Q, v)$*, we have*

$$\sup_{Q \in \mathsf{C}^{\mathrm{E}}(\delta,\iota)} \mathbb{E}_{\gamma\sim Q}\big[\mathrm{Reg}_{i,t}(\gamma)\big] \ , \ \sup_{Q \in \mathsf{C}^{\mathrm{E}}(\delta,\iota)} \mathbb{E}_{\gamma\sim Q}\big[\mathrm{Envy}_{i,t}(\gamma)\big] = \tilde{O}(\sqrt{\iota t} + \sqrt{\delta}\cdot t) \tag{11}$$

and

$$\sup_{Q \in \mathsf{C}^{\mathrm{E}}(\delta,\iota)} \mathbb{E}_{\gamma\sim Q}\left[\|\bar{b}^t - (1/n)1_n\|^2\right] \ , \sup_{Q \in \mathsf{C}^{\mathrm{E}}(\delta,\iota)} \mathbb{E}_{\gamma\sim Q}\left[\|\bar{u}^t - u^*\|^2\right] \ , \sup_{Q \in \mathsf{C}^{\mathrm{E}}(\delta,\iota)} \mathbb{E}_{\gamma\sim Q}\left[\|\bar{u}^t - u^\gamma\|^2\right] = \tilde{O}(\delta + \iota/t) \ .$$

Remark 1 (Markov Input)**.** *We can specialize the result in Theorem 2 to fast mixing or Markov item sequences. Fast mixing means the deviation* δ *decreases exponentially, i.e., for all* $1 \le \iota \le t-1$*,*

$$\sup_{\gamma} \sup_{\tau=1,\ldots,t-\iota} \|Q^{\tau+\iota}(\theta^{1:\tau}) - \Pi\|_{\mathrm{TV}} \le M\rho^\iota \ , \tag{12}$$

for some $M > 0$*,* $\rho \in [0, 1)$*, and* Π *is the stationary distribution. Examples include finite state-space time-homogeneous Markov chains and uniformly ergodic Markov chains on general state spaces (Meyn and Tweedie, 2012, Chapter 16). In these cases, setting* $\iota = \frac{\log t + \log M}{\log(\rho^{-1})} = O(\frac{\log t}{\log(\rho^{-1})})$ *implies* $\delta \le 1/t$*. This means the Markov chain from which* γ *is generated takes* $O(\log t)$ *steps to get* $(1/t)$*-close to stationarity. The dominant term for the regret in Theorem 2 (further ignoring* M*) is then* $(1 + \frac{1}{\log(\rho^{-1})})^{1/2}\sqrt{t}$*. The term in the parenthesis reflects the inflation caused by input dependency. To recover the case of i.i.d. input, we simply send* $\rho \to 0$ *and the usual* $\sqrt{t}$ *regret and envy rates and* $1/t$ *utility and expenditure convergence rates are again recovered.*

3.3 Periodic Input

Item sequences often exhibit statistical periodic structure. For example, when allocating compute time to requestors, there will be more requests during weekdays and less on weekends. The compute request patterns vary throughout the week, and yet the weekly pattern would repeat over time. Similarly, Internet traffic typically exhibits periodic structure.

Formally, assume that the horizon t divides into K blocks of time, each of size q. This divides the item sequence γ into consecutive blocks of length q. Within each block, we allow a distribution with arbitrary dependence between time steps, but we assume that the blocks, as a whole, are identically and independently distributed. We define the set of periodic input distributions as follows:

$$\mathsf{C}^{\mathrm{P}}(q) := \left\{ Q \in \Delta(\Theta^q)^K : Q^{1:q} = Q^{q+1:2q} = \ldots = Q^{t-q+1:t} \right\} . \tag{13}$$

Theorem 3 (Periodic Case)**.** *For the periodic case, Algorithm 1 guarantees that for an instance* $\mathsf{A} = (n, t, \Theta, Q, v)$*, we have*

$$\sup_{Q \in \mathsf{C}^{\mathrm{P}}(q)} \mathbb{E}_{\gamma \sim Q}\left[\mathrm{Reg}_{i,t}(\gamma)\right] , \sup_{Q \in \mathsf{C}^{\mathrm{P}}(q)} \mathbb{E}_{\gamma \sim Q}\left[\mathrm{Envy}_{i,t}(\gamma)\right] = \tilde{O}\left(\sqrt{qt}\right) \tag{14}$$

and

$$\sup_{Q \in \mathsf{C}^{\mathrm{P}}(q)} \mathbb{E}_{\gamma \sim Q}\left[\|\bar{b}^t - (1/n)1_n\|^2\right] , \sup_{Q \in \mathsf{C}^{\mathrm{P}}(q)} \mathbb{E}_{\gamma \sim Q}\left[\|\bar{u}^t - u^*\|^2\right] , \sup_{Q \in \mathsf{C}^{\mathrm{P}}(q)} \mathbb{E}_{\gamma \sim Q}\left[\|\bar{u}^t - u^\gamma\|^2\right] = \tilde{O}(q^2/t) .$$

If the length q of the blocks is of order $o(t)$ then the time-averaged regret and envy are both vanishing. For the i.i.d. case, we can set $q = 1$ to recover the previous results.

Now we comment on dependence on the period length q. Suppose the item sequence consists of K blocks, and blocks are i.i.d. We still allow arbitrary dependence within a block. The proof of Theorem 3 essentially relies on the result (Theorem 9) that DA produces iterates whose squared error is of order $O(q^2/t)$. Consider dual averaging with the knowledge of the block structure q. Then the rate $1/K = q/t$ can be achieved by executing DA using one randomly chosen data point within a block, throwing away the rest in that same block. Such selection produces K i.i.d. samples from the distribution. In comparison, the rate in $O(q^2/t)$ is worse off by a factor of q due to not knowing the block-structure information.

4 Proof Technique: Nonstationary Dual Averaging

The PACE dynamics can be cast as online dual averaging (Xiao, 2010) applied to the dual of the hindsight allocation program in (1). This will be discussed in more detail in Section 4.2 and Appendix C.2. However, in order to characterize its performance under various types of nonstationary input, we first extend the general convergence results for dual averaging to incorporate nonstationary input. The convergence results that we developed for dual averaging under three different types of input models are novel and are of independent interest.

We remark that the results for the stochastic setting given by Xiao (2010) cannot be used directly, since they rely on the stringent i.i.d. assumption. Duchi et al. (2012) consider ergodic mirror descent (MD) for convex problems under some (but not all) of our nonstationary input models. Their results and analysis cannot be used in our case either, since their results do not allow using the composite structure, whose strong convexity we leverage. Moreover, unlike DA, an MD-based approach would require tuning parameters such as stepsizes.

In this section, after introducing the setup of DA in Section 4.1, we present a DA convergence result for independent but not identical input in Section 4.3, for which we outline the proof idea and technical challenges in Section 4.3. Due to space limitations, we present DA convergence results for ergodic and periodic inputs in Appendix D.3.

In the nonstationary setup of DA, we emphasize that whenever we mention convergence, we mean convergence of DA iterates to the population-level optimum w^*_Π (or sometimes the hindsight optimum), up to some error caused by nonstationarity.

4.1 Review: The DA Algorithm

We review the dual averaging setup in the strongly convex case (Xiao, 2010, §1.1). Consider a stochastic optimization problem of the form

$$\min_{w} \left\{ \phi(w) := \mathbb{E}_{z\sim\Pi}\big[F(w,z)\big] = \mathbb{E}_{z\sim\Pi}\big[f(w,z)\big] + \Psi(w) \right\}, \tag{15}$$

where $w \in (\mathbb{R}^d, \|\cdot\|)$ is the variable, Ψ is a closed convex function with closed domain $\operatorname{Dom}\Psi := \{w \in \mathbb{R}^n : \Psi(w) < \infty\}$. The expectation is taken over a probability distribution Π on a measurable space Z. For each $z \in Z$, the function $f(\cdot, z)$ is convex and subdifferentiable (a subgradient always exists) on $\operatorname{Dom}\Psi$. Let $F(w,z) = f(z,w) + \Psi(w)$.

Assumptions. Let $\mathsf{G}(w,z)$ be a fixed element in the set of subgradients $\partial_w f(w,z)$.

1. For almost every z, it holds $\|\mathsf{G}(w,z)\|_* \le G$, where $\|\cdot\|_* = \max_{\|w\|\le 1}\langle s, w\rangle$ is the dual norm.
2. There exists an $\bar{F} \in \mathbb{R}$ such that $F(w,z) \le \bar{F}$ for all $w \in \operatorname{Dom}\Psi$ and (almost every) z.
3. Ψ is σ-strongly convex, i.e., $\Psi(\alpha w + (1-\alpha)u) \le \alpha\Psi(w) + (1-\alpha)\Psi(u) - \frac{\sigma}{2}\alpha(1-\alpha)\|w-u\|^2$ for $w, u \in \operatorname{Dom}\Psi$.

Because of our strong convexity assumption, the solution to (15) is unique. Associated with Π we define $w^*_\Pi := \arg\min \mathbb{E}_{z\sim\Pi}\big[F(w,z)\big]$.

The dual averaging algorithm (DA) (Xiao, 2010, Algorithm 1) aims to produce a sequence converging to the optimal point w^*_Π or minimize the associated regret (Xiao, 2010, §1.2). The algorithmic details for DA are presented in Algorithm 2. Note that although Xiao (2010) only considers the case of i.i.d. data, DA iterates are defined for every input sequence $\{z_\tau\}_{\tau=1}^t$, regardless of any distributional properties of the sequence.

The first step in our analysis is a relationship between regret and the suboptimality $\|w_t - w^*_\Pi\|$ derived by Xiao (2010). Consider the dual averaging algorithm with data $\{z_\tau\}_{\tau=1}^t$. We denote the one-step subgradient by $g_\tau := \mathsf{G}(w_\tau, z_\tau)$ and the average subgradient by $\bar{g}_\tau = (\sum_{s=1}^\tau g_s)/\tau$. Given data $\{z_\tau\}_{\tau=1}^t$, we define the regret and the sum of squared subgradient norms [2]

$$R_t(w) := \sum_{\tau=1}^{t} \big(F(w_\tau, z_\tau) - F(w, z_\tau)\big), \quad \Delta_t := \frac{1}{2\sigma}\big(5\|g_1\|_*^2 + \sum_{\tau=1}^{t-1}\|g_{\tau+1}\|_*^2/\tau\big).$$

The bound $\Delta_t \le (6 + \log t)G^2/(2\sigma)$ holds in a deterministic manner due to the bounded subgradient assumption.

Xiao (2010) shows the following bound on suboptimality of w_t

Lemma 1 (Regret Bound, Section B.2 in Xiao (2010))**.** *For any sequence $\{z_\tau\}_{\tau=1}^t$, any $w \in \operatorname{Dom}\Psi$, any $t = 1, 2, \ldots$, it holds $\|w_{t+1} - w\|^2 \le \frac{2}{\sigma t}\big(\Delta_t - R_t(w)\big)$.*

4.2 Review: PACE as Dual Averaging

In this section we review how to cast PACE as dual averaging applied to the problem (4). This derivation was originally given in Gao et al. (2021). Let $f(\beta,\theta) = \max_i \beta_i v_i(\theta)$, $\Psi(\beta) = -\frac{1}{n}\sum_{i=1}^n \log\beta_i$, in which case we get $F(\beta,\theta) = f(\beta,\theta) + \Psi(\beta) = \max_{i\in[n]}\big\{\beta_i v_i(\theta)\big\} - \frac{1}{n}\sum_{i=1}^n \log\beta_i$. Following (Gao and Kroer, 2021, §5), since $f(\,\cdot\,,\theta)$ is a piecewise linear function, a subgradient is $\mathsf{G}(\beta,\theta) := v_{i^\tau}(\theta)e_{i^\tau} \in \partial_\beta f(\beta,\theta)$, where $i^\tau = \min\{\arg\max_i \beta_i v_i(\theta)\}$ is the index of the winning agent (see, e.g., Beck (2017, Theorem 3.50)). Based on this instantiation of DA, we get that the iterates $\{\beta^\tau\}_{\tau=1}^{t+1}$ generated by the PACE dynamics (Algorithm 1) are exactly the iterates w_τ generated by $\mathrm{DA}(\mathsf{G}, \Psi, \gamma)$ (Algorithm 2). A proof is given in Appendix C.2.

Before moving on to showing our results, let us first touch on the fact that dual averaging provides worst-case regret guarantees. Naively, one may expect that these regret guarantees would directly translate into regret guarantees on the primal performance, meaning a bound on $\sup_\gamma \mathrm{Reg}_{i,t}(\gamma)$.

[2] See the first equation on page 2584 in Xiao (2010). In Xiao (2010)'s notation, set $\beta_\tau = 0$ all $\tau \ge 1$ and $\beta_0 = \sigma$, plug in the bound $h(w_2) \le 2\|g_1\|_*^2/\sigma$ and we have the expression of Δ_t in our paper.

Algorithm 2: $\mathrm{DA}(\mathsf{G}, \Psi, \{z_\tau\}_{\tau=1}^t)$

Input: subgradient G, regularizer Ψ and data $\{z_\tau\}_{\tau=1}^t$.
Initialize: set $\bar{g}_0 = 0$ and $w_1 = \arg\min \Psi$.
for $\tau = 1, \ldots, t$ **do**
  Observe z_τ and compute $g_\tau = \mathsf{G}(w_\tau, z_\tau)$.
  Average subgradients (the *dual average*) via $\bar{g}_\tau = \frac{\tau-1}{\tau}\bar{g}_{\tau-1} + \frac{1}{\tau}g_\tau$.
  Compute the next iterate $w_{\tau+1} = \arg\min_w\{\langle \bar{g}_\tau, w\rangle + \Psi(w)\}$.

However, such dual regret bounds do not imply a worst-case primal regret bound. From a technical perspective, it is unclear how a regret bound on the dual objective would translate to a regret bound on envy or utilities. Secondly, based on results in the online fair allocation literature (Azar et al., 2016; Banerjee et al., 2022), it is known that not only can we not get a no-regret guarantee on the primal performance, it is not even possible to achieve a constant competitive ratio in the worst-case setting.

4.3 Independent Data with Adversarial Corruption

We present a DA convergence result with independent data in this section. Theorem statements for the ergodic case and periodic case are presented in Appendix D.3. Discarding the i.i.d. assumption on the data $\{z_\tau\}_{\tau=1}^t$, we let P be the joint distribution of $\{z_\tau\}_\tau$ and let P^τ be the marginal distribution of z_τ. We study the relationship between the DA iterates w_{t+1} and w_Π^* by bounding the mean-square difference $\mathbb{E}_{\{z_\tau\}_{\tau=1}^t \sim P}\big[\|w_{t+1} - w_\Pi^*\|^2\big]$, and thus demonstrate in what sense the data distribution P should stay close to the i.i.d. distribution Π in order to preserve DA convergence.

We first introduce a variant of $\mathsf{C}^{\mathrm{ID}}(\delta)$ with a target distribution Π:

$$\mathsf{C}^{\mathrm{ID}}(\delta; \Pi) := \left\{ P \in \Delta(\Theta)^t : \frac{1}{t}\sum_{\tau=1}^t \|P^\tau - \Pi\|_{\mathrm{TV}} \le \delta \right\} . \tag{16}$$

Theorem 4 (DA Convergence, Independent Data with Adversarial Corruption)**.** *If $\{z_\tau\}_{\tau=1}^t \sim P$ and $P \in \mathsf{C}^{\mathrm{ID}}(\delta, \Pi)$. Then for $t \ge 1$,*

$$\mathbb{E}_{\{z_\tau\}_{\tau=1}^t \sim P}\big[\|w_{t+1} - w_\Pi^*\|^2\big] \le \frac{(6 + \log t)G^2}{\sigma^2 t} + \frac{8\bar{F}}{\sigma}\delta = \tilde{O}(\delta + 1/t) .$$

Moreover, the rate $\tilde{O}(\delta + 1/t)$ applies to $\mathbb{E}\big[\|w_{t+1} - w_\gamma^*\|_2^2\big]$ and $\mathbb{E}\big[\|w_\gamma^* - w_\Pi^*\|_2^2\big]$ (See Appendix D.1).

Remark 2. *From this result we can tell when DA retains last-iterate convergence. Suppose the number of corrupted data points is of order $o(t)$ (assuming corruption on each data is of the same order), then the corruption per item is $\delta = o(1)$ and DA converges to the optimal solution as $t \to \infty$. Furthermore, if the corruption per data point is of order $\delta = O(1/t)$, then the fast rate $1/t$ is retained.*

By Section 4.2, we get as a corollary that the PACE iterates $\{\beta_t\}$ will converge to β^*, the solution to the infinite-dimensional dual program (4). This is the building block for our results in Section 3.

While the full proof is too long to fit in the paper, we now give a proof sketch to show the main ideas behind how we derive Theorem 4. We begin the proof by noticing Lemma 1 is deterministic and valid for any $\{z_\tau\}_{\tau=1}^t$. Now set $w = w_\Pi^*$ in Lemma 1. If the input data $\{z_\tau\}_{\tau=1}^t$ were i.i.d. from Π, i.e., $P = \Pi^{\otimes t}$, then $\mathbb{E}[R_t(w_\Pi^*)]$ would be greater than zero, and we would obtain

$$\mathbb{E}\big[\|w_{t+1} - w_\Pi^*\|^2\big] \le \frac{2}{\sigma t}\Big(\mathbb{E}[\Delta_t] - \mathbb{E}\big[R_t(w_\Pi^*)\big]\Big) \le \frac{(6 + \log t)G^2}{\sigma^2 t} .$$

However, in the nonstationary case, the regret $\mathbb{E}[R_t(w_\Pi^*)]$ might be negative. At a high level, our results are achieved by introducing appropriate measures of the nonstationarity and then lower bounding $\mathbb{E}[R_t(w_\Pi^*)]$ based on those measures.

To this end, we decompose the regret as follows. Write

$$R_t(w_\Pi^*) = \underbrace{\sum_{\tau=1}^{t} \big(F(w_\tau, z_\tau) - \phi_\Pi(w_\tau)\big) + \sum_{\tau=1}^{t} \big(\phi_\Pi(w_\Pi^*) - F(w_\Pi^*, z_\tau)\big)}_{\mathrm{I}} + \underbrace{\sum_{\tau=1}^{t} \big(\phi_\Pi(w_\tau) - \phi_\Pi(w_\Pi^*)\big)}_{\mathrm{II}} .$$

By optimality of w_Π^* we have $\mathrm{II} \geq 0$. Using the bound on the TV distance between $\{P^\tau\}_\tau$ and Π, and boundedness of F we can control the other two terms. The key is, conditional on $\mathcal{F}_{\tau-1}$, the iterate w_τ is deterministic and the distribution of z_τ is P^τ due to independence assumption. For each term in the first summation, we condition on $\mathcal{F}_{\tau-1}$ and obtain

$$\begin{aligned}
|\mathbb{E}[F(w_\tau, z_\tau) - \phi_\Pi(w_\tau)|\mathcal{F}_{\tau-1}]| &= \left|\mathbb{E}\left[\int_{\mathcal{Z}} F(w_\tau, z) P^\tau(\mathrm{d}z \mid z_{1:\tau-1}) - \int_{\mathcal{Z}} F(w_\tau, z)\, \mathrm{d}\Pi(z) | \mathcal{F}_{\tau-1}\right]\right| \\
&= \left|\mathbb{E}\left[\int_{\mathcal{Z}} F(w_\tau, z) P^\tau(\mathrm{d}z) - \int_{\mathcal{Z}} F(w_\tau, z)\, \mathrm{d}\Pi(z) | \mathcal{F}_{\tau-1}\right]\right| \\
&\leq \mathbb{E}\left[\left|\int_{\mathcal{Z}} F(w_\tau, z) P^\tau(\mathrm{d}z) - \int_{\mathcal{Z}} F(w_\tau, z)\, \mathrm{d}\Pi(z)\right| | \mathcal{F}_{\tau-1}\right] \\
&\leq 2\bar{F} \|P^\tau - \Pi\|_{\mathrm{TV}} .
\end{aligned}$$

The second sum can be handled similarly. For the detailed proof and a generalization see Appendix D.2 in Appendix D.

We now discuss how this paper handles nonstationarity differently from the existing literature. Let us use Balseiro et al. (2020) as a reference point, since they consider very similar categories of nonstationary inputs. From a online constrained optimization perspective, Balseiro et al. (2020) relies on the fact that their objective is separable across timesteps in order to handle nonstationarity. In contrast, this structure does not exist for the online fair allocation problem, because the objective takes the logarithm of the utility over time. Concretely, in Balseiro et al. (2020), the objective is of the form $\sum_{\tau=1}^{t} f_\tau(x)$. This type of time-separability occurs in all works that we are aware of for non-stationary inputs, also e.g. the ergodic mirror descent paper Duchi et al. (2012) (Duchi et al. (2012) is not an online setting, but the expectation in their objective is analogous to time separability). Time separability is used as a key property for deriving regret bounds in those papers. The separability enables translating dual regret to primal regret by a weak duality argument (see e.g. Prop. 1 in Balseiro et al. (2020) where a time-separable dual allows weak duality).

In contrast, our problem has time-separability only in the dual formulation, but not in the primal one, which is where we ultimately want guarantees (since we are interested in utilities converging). Our contribution to handling nonstationarity in online fair allocation is showing that a dual approach works even without the separability structure. We begin by deriving convergence guarantees on the last dual iterate, which we achieve by modifying the dual averaging proof to take into account nonstationarity and analyzing the stability of the dual variables in dual averaging. We note that this technique depends on strong convexity, unlike for the time-separable case. Next, to go from dual last-iterate convergence to primal regret bounds we use the first-order optimality condition for EG program, which are specific techniques for our problem (such techniques were also used in the previous PACE paper Gao et al. (2021)).

5 Conclusion

We established new convergence results for dual averaging under nonstationary data input models, namely, adversarial corruption, ergodic, and block-independent input models. Leveraging these results, we showed that, for online fair allocation problems with item arrivals generated from the above nonstationary data input models, the PACE algorithm automatically adapts to them and achieves asymptotic fairness and efficiency without any parameter tuning. Numerical experiments demonstrated the effectiveness of PACE under these data input models.

Acknowledgments and Disclosure of Funding

This research was supported by the Office of Naval Research Young Investigator Program under grant N00014-22-1-2530, and the National Science Foundation award IIS-2147361.

References

Azar, Y., Buchbinder, N., and Jain, K. (2016). How to allocate goods in an online market? *Algorithmica*, 74(2):589–601.

Balseiro, S., Lu, H., and Mirrokni, V. (2020). The best of many worlds: Dual mirror descent for online allocation problems. *arXiv preprint arXiv:2011.10124*.

Balseiro, S. R. and Gur, Y. (2019). Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 65(9):3952–3968.

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)*, pages 1–19. SIAM.

Bateni, M., Chen, Y., Ciocan, D. F., and Mirrokni, V. (2021). Fair resource allocation in a volatile marketplace. *Operations Research*.

Beck, A. (2017). *First-Order Methods in Optimization*, volume 25. SIAM.

Besbes, O., Gur, Y., and Zeevi, A. (2015). Non-stationary stochastic optimization. *Operations Research*, 63(5):1227–1244.

Birnbaum, B., Devanur, N. R., and Xiao, L. (2011). Distributed algorithms via gradient descent for fisher markets. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 127–136. ACM.

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.

Cheung, Y. K., Cole, R., and Devanur, N. R. (2020). Tatonnement beyond gross substitutes? Gradient descent to the rescue. *Games and Economic Behavior*, 123:295–326.

Cheung, Y. K., Cole, R., and Tao, Y. (2018). Dynamics of distributed updating in fisher markets. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 351–368.

Cole, R., Devanur, N. R., Gkatzelis, V., Jain, K., Mai, T., Vazirani, V. V., and Yazdanbod, S. (2017). Convex program duality, fisher markets, and Nash social welfare. In *18th ACM Conference on Economics and Computation, EC 2017*. Association for Computing Machinery, Inc.

Conitzer, V., Kroer, C., Panigrahi, D., Schrijvers, O., Sodomka, E., Stier-Moses, N. E., and Wilkens, C. (2019). Pacing equilibrium in first-price auction markets. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. ACM.

Conitzer, V., Kroer, C., Sodomka, E., and Stier-Moses, N. E. (2021). Multiplicative pacing equilibria in auction markets. *Operations Research*.

Duchi, J. C., Agarwal, A., Johansson, M., and Jordan, M. I. (2012). Ergodic mirror descent. *SIAM Journal on Optimization*, 22(4):1549–1578.

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

Esfandiari, H., Korula, N., and Mirrokni, V. (2018). Allocation with traffic spikes: Mixing adversarial and stochastic models. *ACM Transactions on Economics and Computation (TEAC)*, 6(3-4):1–23.

Gao, Y. and Kroer, C. (2020). First-order methods for large-scale market equilibrium computation. In *Neural Information Processing Systems 2020, NeurIPS 2020*.

Gao, Y. and Kroer, C. (2021). Infinite-dimension fisher markets and tractable fair division. *arXiv preprint arXiv:2010.03025*.

Gao, Y., Kroer, C., and Peysakhovich, A. (2021). Online market equilibrium with application to fair division. *arXiv preprint arXiv:2103.12936*.

Gkatzelis, V., Psomas, A., and Tan, X. (2021). Fair and efficient online allocations with normalized valuations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5440–5447.

Harper, F. M. and Konstan, J. A. (2016). The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems (TIIS)*, 5(4):19.

Kroer, C., Peysakhovich, A., Sodomka, E., and Stier-Moses, N. E. (2021). Computing large competitive equilibria using abstractions. *Operations Research*.

Kroer, C. and Stier-Moses, N. E. (2022). Market equilibrium models in large-scale internet markets. *Innovative Technology at the Interface of Finance and Operations. Springer Series in Supply Chain Management. Springer Natures, Forthcoming*.

Manshadi, V., Niazadeh, R., and Rodilitz, S. (2021). Fair dynamic rationing. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 694–695.

Meyn, S. P. and Tweedie, R. L. (2012). *Markov Chains and Stochastic Stability*. Springer Science & Business Media.

Nesterov, Y. (2003). *Introductory Lectures on Convex Optimization: A Basic Course*, volume 87. Springer Science & Business Media.

Sinclair, S. R., Banerjee, S., and Yu, C. L. (2021). Sequential fair allocation: Achieving the optimal envy-efficiency tradeoff curve. *arXiv preprint arXiv:2105.05308*.

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

Xiao, L. (2010). Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 11:2543–2596.

Zhang, L. (2011). Proportional response dynamics in the fisher farket. *Theoretical Computer Science*, 412(24):2691–2698.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section **??**.
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] See Section 3
   (b) Did you describe the limitations of your work? [No] Our work is mostly theoretical and studies the fair-allocation algorithm PACE under various nonstationary settings.
   (c) Did you discuss any potential negative societal impacts of your work? [No] Not applicable.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] We read them.
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Section 3 and Section 4.
   (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix C, Appendix D and Appendix E.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See Appendix F.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Appendix F.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] No applicable.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix F.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [No] We use synthetic data.
   (b) Did you mention the license of the assets? [No] N/A.
   (c) Did you include any new assets either in the supplemental material or as a URL? [No] N/A.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No] N/A.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] N/A.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [No] N/A.
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [No] N/A.

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [No] N/A.

A Related Work

Since our work studies competitive equilibrium computation, online fair resource allocation and stochastic optimization, while PACE employs the idea of pacing in auction mechanism design, we further discuss related work in these areas.

Convex optimization for computing competitive equilibria. Convex optimization algorithm (especially first-order methods) and their theory have been used to design and analyze algorithms for computing competitive equilibria, often through equilibrium-capturing convex programs Birnbaum et al. (2011); Cole et al. (2017); Cheung et al. (2020); Gao and Kroer (2020); Gao et al. (2021). Applying a first-order method to such a convex program often leads to (recovers) interpretable market dynamics that emulate real-world economic behaviors, such as the proportional response dynamics Birnbaum et al. (2011); Zhang (2011); Cheung et al. (2018); Gao and Kroer (2020) and tâtonnement (Cheung et al., 2020). The PACE algorithm of Gao et al. (2021) is no exception: it results from applying dual averaging to a specific convex program. Discrete variants of these convex programs have also been used for fair indivisible allocation (Caragiannis et al., 2019), which yields some efficiency and fairness guarantees, though the discreteness breaks the connection to competitive equilibria.

(Online) fair resource allocation. Azar et al. (2016) consider an online Fisher market with arbitrary item arrivals. They focus on a quality measure that is minimized at a competitive equilibrium and give an online algorithm that achieves a competitive ratio logarithmic in the size of the market and the ratio between the maximum and minimum (nonzero) buyer valuations over individual items. This algorithm requires solving a nontrivial linear program per iteration and is not known to improve with stochastic arrivals. Banerjee et al. (2022) considers the problem of online allocation of divisible items to maximize Nash social welfare. They show that, under arbitrary item arrivals but with access to meaningful predictions of each buyer's total utility given all items, an online algorithm of the primal-dual type achieves a logarithmic competitive ratio. Gkatzelis et al. (2021) study the setting where items arrive online and with two agents. They focus on satisfying the no-envy condition while maximizing social welfare, and show that one do this approximately by allocating items proportionally to valuations, assuming that valuations are normalized. Manshadi et al. (2021) studies the problem of rationing a social good and propose simple, implementable algorithms that promote fairness and efficiency. In their setting, it is the agents' demands rather than the supply that are sequentially realized and possibly correlated over time. Bateni et al. (2021) uses Gaussian processes to model item arrivals and consider a budget-weighted proportional fairness metric. They propose a reoptimization policy that consumes buyers' budgets and clears the market gradually while ensuring a competitive ratio in hindsight w.r.t. this metric. This policy periodically resolves the Eisenberg-Gale (EG) convex program and does not require prior knowledge of future item arrivals. Our work differs from the above literature as follows. First, we consider practically-motivated nonstationary data input models for item arrivals that interpolate between fully adversarial and fully stochastic (i.i.d.). Second, we show that the PACE algorithm, without any parameter tuning, adapts to different data input models and achieves strong performance guarantees that depend mildly on the "nonstationarity" of these models. Given that PACE is scalable, interpretable and easy to implement this paper further ensures its effectiveness upon more realistic, non-i.i.d. item arrival processes.

(Nonstationary) stochastic optimization. Many stochastic optimization algorithms have been shown to attain nontrivial performance guarantees under under nonstationary data input (Duchi et al., 2012; Balseiro et al., 2020; Besbes et al., 2015). Motivated by high-dimensional and distributed optimization problems, Duchi et al. (2012) analyzes stochastic mirror descent under ergodic data input. Balseiro et al. (2020) analyzes a version of mirror descent for online resource allocation. They show that it achieve strong regret bounds under different data input models without knowing the model in advance. The ergodic and periodic data input models in this paper are motivated by those considered in Duchi et al. (2012) and Balseiro et al. (2020). Different from these papers which focus on mirror descent, this paper focuses on the dual averaging algorithm, a different stochastic optimization algorithm particularly suitable for the equilibrium-capturing convex program we study. Furthermore,

we achieve stronger results than those past papers, by focusing on a setting where a composite term has strong convexity.

Pacing in auction mechanism design. The PACE algorithm uses first-price auctions with pacing. As noted in Gao and Kroer (2020), the idea of pacing has also been used widely in budget management strategies for Internet advertising auctions, with strong revenue and incentive guarantees (see, e.g., Conitzer et al. (2019, 2021); Balseiro and Gur (2019)). It is also used widely in practice, as reported in Conitzer et al. (2021). As shown in Balseiro et al. (2020), pacing strategies ensure individual bidders' returns on their budgets and, if used by all buyers, lead to approximate Nash equilibria. Similar to the analysis in Gao et al. (2021), in this paper, we focus on competitive equilibrium and fairness properties of PACE, rather than game-theoretic (incentive) properties.

B Review of Linear Fisher Market

In fair division the goal is to perform this allocation in a *fair* way, while simultaneously also guaranteeing some form of efficiency, typically Pareto efficiency. In the case of allocating m divisible goods to n agents, the *competitive equilibrium from equal incomes* (CEEI) allocation guarantees many fairness properties. In CEEI, every agent is endowed with a unit budget of faux currency, a competitive equilibrium is computed, i.e., a set of item prices along with an allocation that clears the market, and the resulting allocation is used as the fair allocation (Varian, 1974). This guarantees several fairness desiderata such as *envy-freeness* (every person prefers their own bundle to that of any other person), *proportionality* (every person prefers their own bundle over receiving their fair share $1/n$ of every item), and Pareto optimality (we cannot make any person better off without making at least one other person worse off).

A linear Fisher market refers to the tuple $\mathsf{F} = (n, m, B, \mathsf{v})$. The market consists of n buyers and m items. We assume each buyer has a budget of B_i. We use $\{1, \ldots, m\}$ to represent the set of items, each of unit supply. The matrix $\mathsf{v} = (\mathsf{v}_1, \ldots, \mathsf{v}_n) \in (\mathbb{R}_+^m)^n$ consists of valuations, with v_i^j being the valuation of item j from buyer i. For buyer i, an allocation of items, $\mathsf{x}_i \in \mathbb{R}_+^m$, gives a utility of $u_i(\mathsf{x}_i) := \langle \mathsf{v}_i, \mathsf{x}_i \rangle := \sum_{j=1}^m \mathsf{v}_i^j \mathsf{x}_i^j$. Note we use different fonts to distinguish notions that appear in both the online allocation problem and the Fisher market.

Definition 1 (Demand)**.** *Given item prices $p \in \mathbb{R}_+^m$, the demand of buyer i is its set of utility-maximizing allocations given the prices and budget:*

$$D_i(p) := \arg\max\{\langle \mathsf{v}_i, \mathsf{x}_i \rangle : \mathsf{x}_i \geq 0,\ \langle p, \mathsf{x}_i \rangle \leq B_i\}\ . \tag{17}$$

Definition 2 (Market Equilibrium)**.** *The market equilibrium of $\mathsf{F} = (n, m, B, \mathsf{v})$ is an allocation-price pair $(\mathsf{x}^*, p^*) \in (\mathbb{R}_+^m)^n \times \mathbb{R}_+^m$ such that the following holds.*

1. *Supply feasibility: $\sum_{i=1}^n \mathsf{x}_i^* \leq 1_m$.*
2. *Buyer optimality: $\mathsf{x}_i^* \in D_i(p^*)$ for all i.*
3. *Market clearance: $\langle p^*, 1_m - \sum_{i=1}^n \mathsf{x}_i^* \rangle = 0$.*

Market equilibrium and fair allocation are related as follows. In CEEI, we construct a mechanism for fair division by giving each agent the same budget of fake currency, i.e., $B_i = B_j$ for all i, j, computing what is called a market equilibrium under this new market, and using the corresponding allocation as our fair allocation rule.

It is known that an allocation x^* from the set of CEEI has many desirable properties. It is Pareto optimal (every market equilibrium is Pareto optimal by the first welfare theorem). It has no envy: since each agent has the same budget in CEEI and every agent is buying something in their demand set, no envy must be satisfied, since they can afford the bundle of any other agent. Finally, proportionality is satisfied, since each agent can afford the bundle where they get $1/n$ of each good.

The ME is essentially a collection of optimization problems (Eq. (17)) coupled through the constraint $\sum_{i=1}^n \mathsf{x}_i \leq \mathbf{1}_m$. A celebrated result is the Eisenberg-Gale convex program, which provides an equivalent characterization of ME.

$$\max_{\mathsf{x}_1, \ldots, \mathsf{x}_n} \sum_{i=1}^n B_i \log \langle \mathsf{v}_i, \mathsf{x}_i \rangle \quad \text{s.t.} \quad \sum_{i=1}^n \mathsf{x}_i^j \leq 1 \ \ \forall j \in [m]\ , \qquad \mathsf{x}_i \in \mathbb{R}_+^m \ \ \forall i \in [n]\ . \tag{18}$$

That is, we maximize the sum of logarithmic utilities under the supply constraint. The solution to the primal problem $\mathsf{x}^* = (\mathsf{x}_1^*, \ldots, \mathsf{x}_m^*)$ along with the vector of dual variables p^* yields a market equilibrium.

The hindsight allocation Eq. (1) is just the EG program of the linear Fisher market $\mathsf{F}_\mathsf{A} = (n, t, B, \mathsf{v})$ where entries of the valuation matrix v are defined by $\mathsf{v}_i^j = v_i(\theta^j)$ for all $i \in [n], j \in [t]$, and $B = (t/n)1_n$.

C PACE as Dual Averaging

In this section, we show how to cast PACE as dual averaging. To this end, we will introduce infinite-dimensional Eisenberg-Gale-type convex programs for the allocation of a (possibly infinite/continuous) set of items. Here, the item supplies correspond to the probability density $\mathrm{d}\bar{Q}/\,\mathrm{d}\mu$ of the average item arrival distribution $\bar{Q}$. They serve as intermediate "reference" convex programs that facilitate the use of DA convergence results developed in the previous section to analyze PACE. When the item space is continuous, the supply function, allocation rules, and the price function in these convex programs are (measurable) functions over such item spaces, which can be infinite-dimensional objects. When item arrivals are drawn from a (fixed) distribution with density s, they correspond to the EG convex programs of the "underlying market" with item supplies s. Note that when the item space Θ is finite, the infinite-dimensional analogues reduce to the classical finite-dimensional EG convex programs. After introducing these concepts we will show that our results on nonstationary DA allows us to derive comparable results on various PACE performance metrics. The results of Gao et al. (2021) cast PACE as dual averaging for the EG convex program of the underlying market, and show guarantees with respect to that program. Here, we will show our results for that setting, as well as for the hindsight allocation problem.

Different from the presentation of DA in Xiao (2010), we do not need an additional regularizer. This paper focuses on DA for strongly convex problems with an existing regularizer in the loss (and hence no auxiliary regularizer is needed), since this is the setting used in the design and analysis of the PACE algorithm; similar convergence results under our new input models can be derived for the general form of DA given in Xiao (2010, Algorithm 1) with an auxiliary regularizer for non-strongly convex problems.

C.1 The Dual of EG and the Infinite-Dimensional Analogue

We derive the dual program of (1). Introduce the dual variables $\beta_i \geq 0$ with $i \in [n]$ for each constraint of the first type and variables $p^\tau \geq 0$ with $\tau \in [t]$ for constraints of the second type. The Lagrangian $L : \mathbb{R}_+^{n\times t} \times \mathbb{R}_+^n \times \mathbb{R}_+^n \times \mathbb{R}_+^t \to \mathbb{R}$ is given by

$$\begin{aligned} L(x, U, \beta, p) &= \sum_{i=1}^n B_i \log U_i + \sum_{i=1}^n \beta_i \big(\langle v_i(\gamma), x_i\rangle - u_i\big) + \sum_{\tau=1}^t p^\tau \left(1 - \sum_{i=1}^n x_i^\tau\right) \\ &= \sum_{\tau=1}^t p^\tau + \sum_{i=1}^n \left(B_i \log U_i - \beta_i u_i\right) + \sum_{i=1}^n \big\langle \beta_i v_i(\gamma) - p, x_i \big\rangle \,. \end{aligned}$$

Maximizing out the variables (x, U) gives the dual program

$$\min_{p\geq 0, \beta \geq 0} \left\{ \sum_{\tau=1}^t p^\tau - \sum_{i=1}^n B_i \log \beta_i + \sum_{i=1}^n \big(B_i \log B_i - B_i\big) \;\middle|\; p \geq \beta_i v_i(\gamma) \;\; \forall i \in [n] \right\} .$$

Dividing by t (recalling $B_i = t/n$), ignoring constants and moving the constraint $p \geq \beta_i v_i(\gamma)$ to the loss, we obtain the following equivalent optimization problem:

$$\min_{\beta \geq 0} \left\{ \frac{1}{t} \sum_{\tau=1}^t \max_{i\in[n]} \beta_i v_i(\theta^\tau) - \frac{1}{n} \sum_{i=1}^n \log \beta_i \right\} . \tag{19}$$

Let β^γ be the optimal solution. To recover the corresponding optimal p^γ we define $p^{\gamma,\tau} = \max_i \beta_i^\gamma v_i(\theta^\tau)$ for $\tau \in [t]$.

We recall the infinite-dimensional analogue of (1):

$$\max_{x \in L_+^\infty(\Theta), u \geq 0} \left\{ \frac{1}{n} \sum_{i=1}^n \log(u_i) \;\middle|\; u_i \leq \langle v_i, x_i \rangle \;\; \forall i \in [n], \;\; \sum_{i=1}^n x_i \leq s \right\}, \tag{20}$$

where $s = \mathrm{d}\bar{Q}/\,\mathrm{d}\mu$ is the average item supply function, and $\langle v_i, x_i \rangle := \int_\Theta v_i x_i \,\mathrm{d}\mu$. The infinite-dimensional analogue of (2) is the following. For any $\delta_0 > 0$,

$$\min_{\beta \geq 0} \left\{ \int_\Theta \Big(\max_{i \in [n]} \beta_i v_i(\theta) \Big) \bar{Q}(\mathrm{d}\theta) - \frac{1}{n} \sum_{i=1}^n \log \beta_i \;\middle|\; \frac{1}{n(1+\delta_0)} \leq \beta_i \leq 1 + \delta_0 \;\; \forall i \in [n] \right\}, \tag{21}$$

A rigorous mathematical treatment of the two infinite-dimensional programs can be found in Gao and Kroer (2021) and (Gao et al., 2021, Section 2). Note the additional constraint in (4) on β does not affect the optimal solution since $1/n \leq \beta_i^* \leq 1$; see Lemma 1 in Gao and Kroer (2021).

The relationship between the finite and infinite versions of (2) is that we have replaced the uniform averaging in (2) with an integral w.r.t. the item average distribution $\bar{Q}$ in (4). For notational simplicity, we suppress dependence on $\bar{Q}$ and let (x^*, u^*, β^*) denote the optimal solutions to the infinite-dimensional programs (3) and (4). Define the corresponding optimal p^* in (4) by $p^* := \max_{i \in [n]} \beta_i^* v_i$.

C.2 PACE as Dual Averaging

In this section we review how to cast PACE as dual averaging applied to the problem (4). This derivation was originally given in Gao et al. (2021). Recall $f(\beta, \theta) = \max_i \beta_i v_i(\theta)$, $\Psi(\beta) = -\frac{1}{n}\sum_{i=1}^n \log \beta_i$ and then

$$F(\beta, \theta) = f(\beta, \theta) + \Psi(\beta) = \max_{i \in [n]} \big\{ \beta_i v_i(\theta) \big\} - \frac{1}{n} \sum_{i=1}^n \log \beta_i \;.$$

Following (Gao and Kroer, 2021, §5), since $f(\,\cdot\,, \theta)$ is a piecewise linear function, a subgradient is

$$\mathsf{G}(\beta, \theta) := v_{i^\tau}(\theta) e_{i^\tau} \in \partial_\beta f(\beta, \theta) \;,$$

where $i^\tau = \min\{\arg\max_i \beta_i v_i(\theta)\}$ is the index of the winning agent (see, e.g., (Beck, 2017, Theorem 3.50)).

Lemma 2 (PACE as Dual Averaging)**.** *The iterates* $\{\beta^\tau\}_{\tau=1}^{t+1}$ *generated by the PACE dynamics (Algorithm 1) are exactly* $\mathrm{DA}(\mathsf{G}, \Psi, \gamma)$ *(Algorithm 2).*

Proof of Lemma 2. Interpret the DA updates using the following substitution: $\Theta \leftrightarrow Z$, $\theta^\tau \leftrightarrow z_\tau$, $\beta^{\tau+1} \leftrightarrow w_{\tau+1}$ and $\bar{g}_{\tau,i} \leftrightarrow \bar{u}_i^\tau$. For initialization in DA, choose β^1 to be the minimizer of Ψ over the cube $[(1+\delta_n)^{-1} n^{-1} 1_n, (1+\delta_0) 1_n]$ and set $\bar{g}_0 = \bar{u}^0 = 0$.

(1) Subgradient computation $\Leftrightarrow$ choose the winning bidder (Line 4 of PACE).

(2) Average subgradient $\Leftrightarrow$ update current averaged utilities (Line 5 of PACE). The i-th entry of $\mathsf{G}(\beta, \theta)$ is exactly the time-τ realized utility of agent i in PACE, that is, $g_{\tau,i} = v_i(\theta^\tau) \mathbb{1}\{i = i^\tau\} = u_i^\tau$. Then the average gradient, $\bar{g}_\tau = \frac{\tau-1}{\tau} \bar{g}_{\tau-1} + \frac{1}{\tau} g_\tau$, is the same as the time-averaged utilities:

$$\bar{g}_{\tau,i} = \frac{\tau - 1}{\tau} \bar{g}_{\tau-1,i} + \frac{1}{\tau} v_i(\theta^\tau) \mathbb{1}\{i = i^\tau\} \;.$$

(3) Solve regularized problem $\Leftrightarrow$ update pacing multiplier (Line 6 of PACE). The minimization problem is separable in agent index i and exhibits a simple and explicit solution. Recall $\bar{g}_{\tau,i} = \bar{u}_i^\tau$:

$$\beta_i^{\tau+1} = \arg\min \left\{ \bar{g}_{\tau,i} \beta_i - \frac{1}{n} \log \beta_i \Big| \frac{1}{n(1+\delta_0)} \leq \beta_i \leq 1 + \delta_0 \right\} \;\Rightarrow\; \beta_i^{\tau+1} = \Pi_{[\ell,h]} \left(\frac{1}{n \bar{u}_i^\tau} \right) \;.$$

□

C.3 Performance Guarantees via Dual Averaging

Now that we have cast PACE as an instantiation of dual averaging and developed results for convergence in nonstationary settings, the following theorems follow easily from the general convergence results for DA. Recall the hindsight optimum β^γ is defined in (2), and its infinite-dimensional counterpart β^* is defined in (4).

Theorem 5 (Convergence of PACE, the Independent Case)**.** *Assume the item sequence $\gamma \sim Q$ and $Q \in \mathsf{C}^{\mathrm{ID}}(\delta)$. Choose $\delta_0 = 1$ in PACE. It holds for $t \geq 1$,*

$$\mathbb{E}\big[\|\beta^t - \beta^*\|^2\big] \leq \frac{(6+\log t)n^2|v|_\infty^2}{t} + 8n|v|_\infty \cdot \delta = \tilde{O}(\delta + 1/t) \ . \tag{22}$$

Moreover, the rate $\tilde{O}(\delta + 1/t)$ applies to $\mathbb{E}[\|\beta^\gamma - \beta^*\|^2]$ and $\mathbb{E}[\|\beta^{t+1} - \beta^\gamma\|^2]$.

Proof. Set $P = Q$, $\Pi = \bar{Q}$, $\sigma = 1/n$, and $\bar{F} = |v|_\infty$ in Theorem 4. □

Here the convergence of β^t to the hindsight counterpart β^γ is of practical importance. This is because β^γ can always be computed after the fact, while its infinite-dimensional counterpart β^* is not necessarily obtainable.

Theorem 6 (Convergence of PACE, the Ergodic and the Periodic Cases)**.** *For the erogdic case, i.e., $\gamma \sim Q$ and $Q \in \mathsf{C}^{\mathrm{E}}(\delta, \iota)$, it holds for $t \geq 1$,*

$$\mathbb{E}[\|\beta^{t+1} - \beta^*\|^2] \leq \frac{C_{E,1} + C_{E,2} \cdot \iota}{t} + C_{E,3} \cdot \delta = \tilde{O}\left(\delta + \frac{\iota}{t}\right) \ .$$

where $C_{E,1} = n^2|v|_\infty^2\big(6 + \log t\big)$, $C_{E,2} = 4n\big(|v|_\infty^2(1+\log t) + |v|_\infty\big)$ and $C_{E,3} = 8n|v|_\infty$.

For the periodic case, i.e., $\gamma \sim Q$ and $Q \in \mathsf{C}^{\mathrm{P}}(q)$, it holds for $t \geq 1$

$$\mathbb{E}[\|\beta^{t+1} - \beta^*\|^2] \leq \frac{C_{P,1} + C_{P,2} \cdot q^2}{t} = \tilde{O}(q^2/t) \ .$$

where $C_{P,1} = C_{E,1}$ and $C_{P,2} = 2n|v|_\infty^2(1+\log t)$.

For both cases, similar convergence results can be stated for $\mathbb{E}[\|\beta^\gamma - \beta^*\|^2]$ and $\mathbb{E}[\|\beta^{t+1} - \beta^\gamma\|^2]$ and are omitted here.

Proof. For the first inequality, set $P = Q$, $\Pi = \bar{Q}$, $\sigma = 1/n$, and $\bar{F} = |v|_\infty$ in Theorem 8. For the second inequality, set additionally $\delta^{\mathrm{b}} = 0$ and $|\mathcal{P}|_\infty = q$ in Theorem 9. □

C.4 From Dual EG Performance Bounds to Primal Performance Bounds

Convergence of β^τ to β^* implies the convergence of the average utilities and expenditure to their infinite-dimensional counterparts. This follows almost directly from results developed by Gao et al. (2021). In particular, they show:

Lemma 3 (PACE Long-Run Behavior, Gao et al. (2021))**.** *For any distribution $Q \in \Delta(\Theta^t)$, let $\gamma \sim Q$. It holds for $t \geq 1$,*

$$\mathbb{E}\big[\|\bar{b}^t - (1/n)1_n\|^2\big] \leq 2\mathbb{E}[\|\beta^{t+1} - \beta^*\|^2] + 4|v|_\infty^2\left(\frac{1}{t}\sum_{\tau=1}^{t}\mathbb{E}[\|\beta^\tau - \beta^*\|^2]\right) ,$$

and

$$\mathbb{E}\big[\|\bar{u}^t - u^*\|^2\big] \leq C_u \cdot \mathbb{E}[\|\beta^{t+1} - \beta^*\|^2] \ ,$$

where $C_u = n^2\left(|v|_\infty^2/\delta_0^2 + (1+\delta_0)^2\right)$.

Finally, we relate results in Appendix C.3 to the main quantities of interest: regret and envy, defined in (6) and (7), as well as convergence to the hindsight utilities. Similar results are given by Gao et al. (2021), and our proof is almost identical to theirs, simply extended to the nonstationary case as well as to the hindsight allocation problem.

Lemma 4 (Regret and Envy). *For any distribution $Q \in \Delta(\Theta^t)$, let $\gamma \sim Q$. It holds for $t \geq 1$,*

$$\mathbb{E}[\|\bar{u}^t - u^\gamma\|^2] \leq C_{r,1} \cdot \mathbb{E}[\|\beta^{t+1} - \beta^*\|^2] + n \cdot C_{r,2} \cdot \left(\frac{1}{t} \sum_{\tau=1}^{t} \mathbb{E}[\|\beta^\tau - \beta^*\|^2] \right) ,$$

$$\mathbb{E}\big[\mathrm{Reg}_{i,t}(\gamma)\big] \leq t \cdot \sqrt{C_{r,1} \cdot \mathbb{E}[\|\beta^{t+1} - \beta^*\|^2] + C_{r,2} \cdot \left(\frac{1}{t} \sum_{\tau=1}^{t} \mathbb{E}[\|\beta^\tau - \beta^*\|^2] \right)} ,$$

and

$$\mathbb{E}\big[\mathrm{Envy}_{i,t}(\gamma)\big] \leq t \cdot \sqrt{C_{e,1} \cdot \mathbb{E}[\|\beta^{t+1} - \beta^*\|^2] + C_{e,2} \cdot \left(\frac{1}{t} \sum_{\tau=1}^{t} \mathbb{E}[\|\beta^\tau - \beta^*\|^2] \right)} ,$$

where $C_{r,1} = 2C_u$, $C_{r,2} = 2n^2|v|_\infty^2$, $C_{e,1} = 2(1+n^2)C_u$ and $C_{e,2} = 4|v|_\infty^2 n^2 + 2n^3$.

Now we have all the ingredients to prove the convergence of PACE.

Proof of Theorem 1, Theorem 2 and Theorem 3. Combine Lemma 3 with Theorem 5 and Theorem 6 and we obtain the first set of inequalities in Theorems 1 to 3. Then combine Lemma 4 with Theorem 5 and Theorem 6, and we obtain the second set of inequalities.

□

D Proofs for Nonstationary Dual Averaging

The DA algorithm Algorithm 2 is obtained in Xiao (2010, §3.2) for the strongly convex case. We simply set $h = (1/\sigma)\Psi$, $\beta_0 = \sigma$ and $\beta_t = 0$ for $t \geq 1$ in (Xiao, 2010, Algorithm 1).

Recall P^τ is the distribution of z_τ. For integers τ and τ' ($\tau' \geq \tau$), let $P^{\tau'}(\cdot \mid z_{1:\tau})$ denote the distribution of $z_{\tau'}$ if the process starts at $z_{1:\tau} = \{z_1, \ldots, z_\tau\}$.

D.1 Remark: Convergence to the Hindsight Optimum

Before developing the nonstationary convergence theory, we digress a bit and introduce a simple deduction through which we can easily show the convergence of DA iterates to the optimum of the hindsight problem based on convergence to w_Π^*. Given data $\{z_\tau\}_{\tau=1}^t$, define the sum $\phi_\gamma(w) = (1/t) \cdot \sum_{\tau=1}^t F(w, z_\tau)$ and its unique minimizer $w_\gamma^* = \arg\min \phi_\gamma(w)$. We claim all results developed for $\|w_{t+1} - w_\Pi^*\|^2$ will also hold for the hindsight suboptimality $\|w_{t+1} - w_\gamma^*\|^2$.

Note the following inequality:

$$R_t(w_\gamma^*) = \sum_{\tau=1}^{t} \big(F(w_\tau, z_\tau) - F(w_\gamma^*, z_\tau)\big) = \sum_{\tau=1}^{t} \big(F(w_\tau, z_\tau) - \phi_\gamma(w_\gamma^*)\big) \geq \sum_{\tau=1}^{t} \big(F(w_\tau, z_\tau) - \phi_\gamma(w_\Pi^*)\big) ,$$

the last term being exactly $R_t(w_\Pi^*)$. Choose $w = w_\gamma^*$ in Lemma 1 and we obtain

$$\mathbb{E}\big[\|w_{t+1} - w_\gamma^*\|_2^2\big] \leq \frac{1}{\sigma t}\Big(\mathbb{E}[\Delta_t] - \mathbb{E}\big[R_t(w_\gamma^*)\big]\Big) \leq \frac{1}{\sigma t}\Big(\mathbb{E}[\Delta_t] - \mathbb{E}\big[R_t(w_\Pi^*)\big]\Big) .$$

It follows that all lower bounds for the regret $\mathbb{E}\big[R_t(w_\Pi^*)\big]$ can be turned into an upper bound for the hindsight suboptimality measure $\|w_{t+1} - w_\Pi^*\|^2$. Convergence to the hindsight optimum is of practical importance since the hindsight optimum w_γ^* can typically be computed, where this is not always the case for the population optimum w_Π^*.

A simple consequence of the deduction above is the following. We note the inequality

$$\|w_\gamma^* - w_\Pi^*\|^2 \leq 2\|w_{t+1} - w_\Pi^*\|^2 + 2\|w_{t+1} - w_\gamma^*\|^2 .$$

Therefore convergence results for $\|w_{t+1} - w_\Pi^*\|$ and $\|w_{t+1} - w_\gamma^*\|^2$ similarly hold for $\|w_\gamma^* - w_\Pi^*\|^2$.

D.2 Proof for Adversarial Corruption and Independent Data

Assume the data $\gamma = \{z_\tau\}_{\tau=1}^t$ follows the distribution P with no further assumptions. We let $z_{1:0} = \emptyset$ and further let $P^\tau(\cdot \mid z_{1:0}) = P^\tau$, the marginal distribution of z_τ.

Define the progressive deviation from Π

$$\delta_t := \sup_{z_1,\ldots,z_t} \sum_{\tau=1}^t \|P^\tau(\cdot \mid z_{1:\tau-1}) - \Pi\|_{\mathrm{TV}} \ . \tag{23}$$

If the data are independent, then it holds

$$\delta_t = \sum_{\tau=1}^t \|P^\tau - \Pi\|_{\mathrm{TV}} \ .$$

Note for the independent case, $\delta_t = t \cdot \delta$ where δ is in Eq. (16). If the data further has identical distribution Π then $\delta_t = 0$. If $\delta_t = O(\log t)$ we call data has mild corruption.

Theorem 7 (Corrupted Data, Generalization of Theorem 4)**.** *It holds*

$$\mathbb{E}\big[\|w_{t+1} - w_\Pi^*\|_2^2\big] = O\Big(\frac{\log t}{\sigma^2 t} + \frac{\delta_t}{\sigma t}\Big) \ , \tag{24}$$

where O hides dependence on constants, G and $\bar{F}$. Recall σ is the strong convexity parameter of Ψ.

Remark 3. *In either the i.i.d. case or the mild corruption case ($\delta_t = O(\log t)$), we recover the usual $O(\log t/t)$ rate.*

Proof of Theorem 7 In Lemma 1, the term Δ_t is upper bounded in a deterministic manner. So it remains to handle R_t. In the i.i.d. case, $\mathbb{E}R_t(w_\Pi^*)$ is positive and thus can be dropped:

$$\mathbb{E}[R_t(w_\Pi^*)] = \mathbb{E}\Big[\sum_{\tau=1}^t \big(F(w_\tau, z_\tau) - F(w_\Pi^*, z_\tau)\big)\Big] = \sum_{\tau=1}^t \big(\phi_\Pi(w_\tau) - \phi_\Pi(w_\Pi^*)\big) \geq 0 \ .$$

However, to handle corrupted data, we need to use

Lemma 5. *The regret can be lower bounded by the corruption parameter δ:*

$$\mathbb{E}[R_t(w_\Pi^*)] \geq -4 \cdot \bar{F}\delta_t \ .$$

Plugging in the above lemma, we get

$$\mathbb{E}[\|w_{t+1} - w\|^2] \leq \frac{1}{\sigma t}\big(\mathbb{E}[\Delta_t] - \mathbb{E}[R_t(w_\Pi^*)]\big) \leq \frac{\big(6 + \log t\big)G^2}{\sigma} + \frac{8\bar{F}}{\sigma}\delta = O\Big(\frac{G^2 \log t}{\sigma^2 t} + \frac{\bar{F}\delta_t}{\sigma t}\Big) \ . \tag{25}$$

Thus, to complete the proof of Theorem 7 we only need to prove Lemma 5.

Proof of Lemma 5. Write

$$\begin{aligned} R_t(w_\Pi^*) = &\sum_{\tau=1}^t \big(F(w_\tau, z_\tau) - \phi_\Pi(w_\tau)\big) && \text{(I)} \\ &+ \sum_{\tau=1}^t \big(\phi_\Pi(w_\Pi^*) - F(w_\Pi^*, z_\tau)\big) && \text{(II)} \\ &+ \sum_{\tau=1}^t \big(\phi_\Pi(w_\tau) - \phi_\Pi(w_\Pi^*)\big) \ . && \text{(III)} \end{aligned}$$

By optimality of w_Π^* we have $\mathrm{III} \geq 0$.

<u>Bounding I and II.</u> Conditional on $\mathcal{F}_\tau$, the iterate w_τ is deterministic and the distribution of z_τ is $P^\tau(\cdot \mid z_{1:\tau-1})$.

Note

$$\mathbb{E}[F(w_\tau, z_\tau) - \phi_\Pi(w_\tau)] = \mathbb{E}[\mathbb{E}[F(w_\tau, z_\tau) - \phi_\Pi(w_\tau)|\mathcal{F}_{\tau-1}]] \ .$$

Let us investigate the inner expectation. Conditional on $\mathcal{F}_{\tau-1}$, the iterate w_τ is deterministic, and the distribution of $z_\tau|\mathcal{F}_{\tau-1}$ is $P^\tau(\cdot \mid z_{1:\tau-1})$ by definition.

$$\begin{aligned}
|\mathbb{E}[F(w_\tau, z_\tau) - \phi_\Pi(w_\tau)|\mathcal{F}_{\tau-1}]| &= \left|\mathbb{E}\left[\int_{\mathcal{Z}} F(w_\tau, z) P^\tau(\mathrm{d}z \mid z_{1:\tau-1}) - \int_{\mathcal{Z}} F(w_\tau, z)\,\mathrm{d}\Pi(z)|\mathcal{F}_{\tau-1}\right]\right| \\
&\leq \mathbb{E}\left[\left|\int_{\mathcal{Z}} F(w_\tau, z) P^\tau(\mathrm{d}z \mid z_{1:\tau-1}) - \int_{\mathcal{Z}} F(w_\tau, z)\,\mathrm{d}\Pi(z)\right||\mathcal{F}_{\tau-1}\right] \\
&\leq \bar{F}\int_{\mathcal{Z}} |\,\mathrm{d}P^\tau(\cdot \mid z_{1:\tau-1}) - \mathrm{d}\Pi(z)| \\
&= 2\bar{F}\cdot\|P^\tau(\cdot \mid z_{1:\tau-1}) - \Pi\|_{\mathrm{TV}} \ .
\end{aligned}$$

where we use boundedness of F, i.e., $\sup_w F(w, z) \leq \bar{F}$ for Π-almost every z,

Next, sum over $\tau = 1, \ldots, t$ and move $|\cdot|$ inside the sum and the outer expectation.

$$\begin{aligned}
|\mathbb{E}[\mathrm{I}]| &= \left|\sum_{\tau=1}^{t}\mathbb{E}[\mathbb{E}[F(w_\tau, z_\tau) - \phi_\Pi(w_\tau)|\mathcal{F}_{\tau-1}]]\right| \\
&\leq \sum_{\tau=1}^{t}\mathbb{E}\Big[\big|\mathbb{E}[F(w_\tau, z_\tau) - \phi_\Pi(w_\tau)|\mathcal{F}_{\tau-1}]\big|\Big] \\
&\leq 2\bar{F}\cdot\sum_{\tau=1}^{t}\mathbb{E}\big[\|P^\tau(\cdot \mid z_{1:\tau-1}) - \Pi\|_{\mathrm{TV}}\big] \\
&\leq 2\bar{F}\cdot\sup_{z_1,\ldots,z_t}\sum_{\tau=1}^{t}\|P^\tau(\cdot \mid z_{1:\tau-1}) - \Pi\|_{\mathrm{TV}} \\
&= 2\bar{F}\delta_t \ .
\end{aligned}$$

Next consider $|\mathbb{E}[\mathrm{II}]|$. The analysis goes through without the outer expectation.

Combining we get

$$\mathbb{E}R_t(w_\Pi^*) = \mathbb{E}[\mathrm{I} + \mathrm{II} + \mathrm{III}] \geq \mathbb{E}[\mathrm{I} + \mathrm{II}] \geq -\big(|\mathbb{E}[\mathrm{I}]| + |\mathbb{E}[\mathrm{II}]|\big) \geq -4\bar{F}\delta_t \ .$$

This completes the proof of Lemma 5. □

D.3 Theorem Statements for Markov Case and Periodic Case

Results for other input types, $\mathsf{C}^{\mathrm{E}}(\delta, \iota)$ and $\mathsf{C}^{\mathrm{P}}(q)$, can be obtained by using more complicated regret decompositions and conditioning arguments. We state the resulting convergence results here. The proofs can be found in Appendix D.4 and Appendix D.5.

Theorem 8 (DA Convergence, Ergodic Case)**.** *For the input distribution P define the ι-step deviation from Π for an integer $1 \leq \iota \leq t-1$:*

$$\epsilon_t(\iota) := \sup_{z_1,\ldots,z_t}\sup_{\tau=1,\ldots,t-\iota}\|P^{\tau+\iota}(\cdot \mid z_{1:\tau}) - \Pi\|_{\mathrm{TV}} \ .$$

Then, for all $t \geq 1$ and any $1 \leq \iota \leq t-1$,

$$\begin{aligned}
\mathbb{E}_{\{z_\tau\}_{\tau=1}^t \sim P}\big[\|w_{t+1} - w_\Pi^*\|_2^2\big] &\leq \frac{\big(6 + \log t\big)G^2}{\sigma^2 t} + \frac{2\big(4\bar{F}\epsilon_t(\iota)t + 2G^2\iota(\log t + 1) + 2\iota\bar{F}\big)}{\sigma t} \\
&= \tilde{O}(\epsilon_t(\iota) + \iota/t) \ .
\end{aligned}$$

Moreover, the rate $\tilde{O}(\epsilon_t(\iota) + \iota/t)$ applies to $\mathbb{E}\big[\|w_{t+1} - w_\gamma^*\|_2^2\big]$ and $\mathbb{E}\big[\|w_\gamma^* - w_\Pi^*\|_2^2\big]$ by the deduction in Appendix D.1.

Remark 4 (Comparison with EMD Duchi et al. (2012))**.** *Now we specialize Theorem 8 to the setting of Remark 1, and we briefly compare our result with the Ergodic Mirror Descent (EMD) results of Duchi et al. (2012). EMD considers nonsmooth convex optimization problems of the form* $f^* = \min\{f(w) = \mathbb{E}_\Pi[F(w;\xi)] \,|\, w \in \mathcal{W}\}$ *for a closed convex set* $\mathcal{W}$*. Differently from our setting, they do not assume strong convexity in* f*, and do not allow a composite term* Ψ *which is not linearized. Assume the Markov chain that generates* $\{z_\tau\}_{\tau=1}^t$ *are fast mixing with* $\epsilon_t(\iota) \le M\rho^\iota$ *for some* $M > 0$ *and* $\rho \in [0,1)$*, then the EMD algorithm produces iterates that satisfy the following convergence rate* [3]

$$\mathbb{E}\big[f(w_{t+1}) - f^*\big] = \tilde{O}\left(\left(1 + \frac{1}{\log(\rho^{-1})}\right)^{1/2} \cdot \frac{1}{\sqrt{t}}\right).$$

As in Remark 1, for the same fast mixing Markov chain, we set $\iota = O\left(\frac{\log t}{\log(\rho^{-1})}\right)$ *and* $\epsilon_t(\iota) = 1/t$ *in Theorem 8, we obtain the rate*

$$\mathbb{E}\big[\|w_{t+1} - w^*\|^2\big] = \tilde{O}\left(\left(1 + \frac{1}{\log(\rho^{-1})}\right) \cdot \frac{1}{t}\right).$$

which is also the rate for $\mathbb{E}[f(w_{t+1}) - f^*]$*. Both results characterize the dependence of convergence rate on* ρ*, the mixing parameter of the Markov chain. However, our result exploits the strong convexity of the optimization problem and achieves the faster rate* $1/t$*, while also achieving convergence in iterates rather than only in function values.*

Theorem 9 (DA Convergence, Block-Independent Case)**.** *Fix an integer* $K \ge 1$*. Let* $\{1 = \tau_1 < \tau_2 < \ldots \tau_{K+1} = t\}$ *be an increasing subsequence of* $[t]$*. Using each two consecutive points, form the interval* $I_k := [\tau_k, \tau_{k+1} - 1]$*. Then* $\mathcal{P} := \{I_k\}_{k=1}^K$ *is a partition of* $[t]$*. Define by* $|I_k| := \tau_{k+1} - \tau_k \ge 1$ *the length of the interval and* $|\mathcal{P}|_\infty := \max_k |I_k|$ *the maximum length of the intervals. Associated with the input distribution* P *and the partition* $\mathcal{P}$ *define the block-wise deviation from* Π *by*

$$\delta^{\mathrm{b}} := \frac{1}{t}\sum_{k=1}^K |I_k| \cdot \left\|\Pi - \frac{1}{|I_k|}\sum_{\tau \in I_k} P^\tau\right\|_{\mathrm{TV}}.$$

Assume $\{z_\tau\}_{\tau=1}^t$ *are block-wise independent according to the partition* $\mathcal{P}$*. Then, for all* $t \ge 1$*,*

$$\begin{aligned}\mathbb{E}_{\{z_\tau\}_{\tau=1}^t \sim P}\big[\|w_{t+1} - w_\Pi^*\|_2^2\big] &\le \frac{(6 + \log t)G^2}{\sigma^2 t} + \frac{2\big(4\bar{F} \cdot \delta^{\mathrm{b}} t + G^2 |\mathcal{P}|_\infty^2 (\log t + 1)\big)}{\sigma t} \\ &= \tilde{O}(\delta^{\mathrm{b}} + |\mathcal{P}|_\infty^2/t)\,. \end{aligned} \tag{26}$$

Moreover, the rate $\tilde{O}(\delta^{\mathrm{b}} + |\mathcal{P}|_\infty^2/t)$ *applies to* $\mathbb{E}\big[\|w_{t+1} - w_\gamma^*\|_2^2\big]$ *and* $\mathbb{E}\big[\|w_\gamma^* - w_\Pi^*\|_2^2\big]$ *by the deduction in Appendix D.1.*

Let us briefly we comment on the dependence on $|P|_\infty$. Suppose there are in total K blocks, each of equal length $|\mathcal{P}|_\infty = q$, and blocks are i.i.d. We still allow arbitrary dependence within a block. Moreover, we choose $\Pi = \bar{P}$ in the definition of δ^{b}. This implies $\delta^{\mathrm{b}} = 0$ and then the rate in Theorem 9 is q^2/t.

Consider dual averaging with the knowledge of the block structure q. Then the rate $1/K = q/t$ can be achieved by executing DA using one randomly chosen data point within a block, throwing away the rest in that same block. Such selection produces K i.i.d. samples from $\bar{P}$. In comparison, the rate in Eq. (26) is worse off by a factor of q due to not knowing the block-structure information.

D.4 Proofs for Markov Case

Now we consider data that are not necessarily independent across time. We restrict our attention to ergodic processes, meaning data tend to be independent as they grow apart in time.

[3] In Eq. (3.2) of Duchi et al. (2012), set $\kappa_1 = M$ and $\kappa_2 = 1/\log(\rho^{-1})$ and ignore the parameters (G, D, κ_1).

Define the ι-step deviation from stationarity

$$\epsilon_t(\iota) := \sup_{z_1,\ldots,z_t} \sup_{\tau=1,\ldots,t-\iota} \|P^{\tau+\iota}(\cdot \mid z_{1:\tau}) - \Pi\|_{\mathrm{TV}} \ .$$

An equivalent quantity is the ϵ-mixing time Duchi et al. (2012)

$$\iota_{\mathrm{mix}}(\epsilon) := \min\Big\{\iota : 1 \le \iota \le t-1, \sup_{z_1,\ldots,z_t} \sup_{\tau=1,\ldots,t-\iota} \|P^{\tau+\iota}(\cdot \mid z_{1:\tau}) - \Pi\|_{\mathrm{TV}} \le \epsilon\Big\} \ .$$

This means, no matter where and when we start the process, it takes only ι steps to get $\epsilon_t(\iota)$-close to the stationary distribution Π. One could expect the deviation $\epsilon_t(\iota)$ decreases as ι increases. This makes sense because for large ι, the process has run long enough to reach stationarity.

Theorem 10 (Mixing Data, Restatement of Theorem 8)**.** *It holds for all $t \ge 1$ and any $1 \le \iota \le t-1$,*

$$\mathbb{E}\big[\|w_{t+1} - w_\Pi^*\|_2^2\big] = O\Big(\frac{\log t}{\sigma^2 t} + \frac{\iota \log t}{\sigma t} + \epsilon_t(\iota)/\sigma\Big) \ , \tag{27}$$

where $O(\cdot)$ hides dependence on constants, G and $\bar{F}$. Here there is a trade-off in ι in the last two terms.

Proof of Theorem 10 We use the proof in Duchi et al. (2012); see Eq. (6.2) in the paper. Decompose $R_t(w_\Pi^*)$ as follows.

$$\begin{aligned}
R_t(w_\Pi^*) = &\sum_{\tau=1}^{t-\iota}\Big(\big(F(w_\tau, z_{\tau+\iota}) - F(w_\Pi^*, z_{\tau+\iota})\big) - \big(\phi_\Pi(w_\tau) - \phi_\Pi(w_\Pi^*)\big)\Big) && \text{(A)}\\
&+ \sum_{\tau=1}^{t-\iota}\big(F(w_{\tau+\iota}, z_{\tau+\iota}) - F(w_\tau, z_{\tau+\iota})\big) && \text{(B)}\\
&+ \sum_{\tau=1}^{t-\iota}\big(\phi_\Pi(w_\tau) - \phi_\Pi(w_\Pi^*)\big) && \text{(C)}\\
&+ \sum_{\tau=1}^{\iota}\big(F(w_\tau, z_\tau) - F(w_\Pi^*, z_\tau)\big) \ . && \text{(D)}
\end{aligned}$$

By optimality of w_Π^* we have $\mathrm{C} \ge 0$. By boundedness of F we get $|\mathrm{D}| \le 2\iota\bar{F}$. Remains to handle A and B. We will show

$$\text{" } \mathrm{A} \le \epsilon_t(\iota)t, \quad \mathrm{B} \le \iota \text{ "} .$$

<u>Bounding A.</u> The key is $z_{\tau+\iota}$ is almost independent of $\mathcal{F}_{\tau-1}$ if ι is moderately large. For each τ,

$$\begin{aligned}
&|\mathbb{E}[F(w_\tau, z_{\tau+\iota}) - \phi_\Pi(w_\tau)]|\\
&= |\mathbb{E}[\mathbb{E}[F(w_\tau, z_{\tau+\iota}) - \phi_\Pi(w_\tau)|\mathcal{F}_{\tau-1}]]|\\
&= \left|\mathbb{E}\left[\mathbb{E}\left[\int_{\mathcal{Z}} F(w_\tau, z)P^{\tau+\iota}(\mathrm{d}z \mid z_{1:\tau-1}) - \int_{\mathcal{Z}} F(w_\tau, z)\Pi(\mathrm{d}z)|\mathcal{F}_{\tau-1}\right]\right]\right| && \text{(Key)}\\
&\le \mathbb{E}\left[\mathbb{E}\left[\left|\int_{\mathcal{Z}} F(w_\tau, z)P^{\tau+\iota}(\mathrm{d}z \mid z_{1:\tau-1}) - \int_{\mathcal{Z}} F(w_\tau, z)\Pi(\mathrm{d}z)\right||\mathcal{F}_{\tau-1}\right]\right]\\
&\le 2\bar{F} \cdot \mathbb{E}\big[\|P^{\tau+\iota}(\cdot \mid z_{1:\tau-1}) - \Pi\|_{\mathrm{TV}}\big]\\
&\le 2\bar{F} \cdot \sup_{z_1,\ldots,z_{\tau-1}} \|P^{\tau+\iota}(\cdot \mid z_{1:\tau-1}) - \Pi\|_{\mathrm{TV}} \le 2\bar{F}\epsilon_t(\iota) \ .
\end{aligned}$$

Analysis for $|\mathbb{E}[F(w_\Pi^*, z_{\tau+\iota}) - \phi_\Pi(w_\Pi^*)]|$ is almost identical. Next sum over $\tau = 1, \ldots, t-\iota$.

$$|\mathbb{E}[\mathrm{A}]| \le 4\bar{F}\epsilon_t(\iota) \cdot t \ .$$

<u>Bounding B.</u> The change in F by ι steps of updates, starting from w_τ, is controlled by $c \cdot \iota G \cdot \frac{1}{\tau}$ where $1/\tau$ acting like a stepsize.

Lemma 6. *Let* $\Pi_{\Psi,\mathcal{W}}(g) := \arg\min_{w\in\mathcal{W}}\{\langle g, w\rangle + \Psi(w)\}$. *If* Ψ *is* σ*-strongly convex, then*

$$\|\Pi_{\Psi,\mathcal{W}}(g) - \Pi_{\Psi,\mathcal{W}}(g')\| \le (1/\sigma)\|g - g'\|_* \ .$$

Proof. See (Nesterov, 2003, Lemma 6.1.2). □

Noting $w_{\tau+1} = \Pi_{\Psi,\mathcal{W}}(\bar{g}_\tau)$ and $w_\tau = \Pi_{\Psi,\mathcal{W}}(\bar{g}_{\tau-1})$, Lemma 6 gives

$$\|w_{\tau+1} - w_\tau\| \le \|\bar{g}_\tau - \bar{g}_{\tau-1}\|_*/\sigma = \|\bar{g}_{\tau-1} - g_\tau\|_*/(\tau\sigma) \le 2G/(\tau\sigma) \ . \tag{28}$$

It holds Π-a.s. that for each τ, the map $w \mapsto F(w, z_{\tau+\iota})$ is Lipschitz with parameter G.

$$\begin{aligned}
|\mathbb{E}[F(w_{\tau+\iota}, z_{\tau+\iota}) - F(w_\tau, z_{\tau+\iota})]| &\le G \cdot \mathbb{E}\big[\|w_{\tau+\iota} - w_\tau\|\big] \\
&\le G \cdot \sum_{t'=\tau}^{\tau+\iota-1} \mathbb{E}\big[\|w_{t'+1} - w_{t'}\|\big] \\
&\le G \cdot \sum_{t'=\tau}^{\tau+\iota-1} 2G/(\sigma t') \\
&\le G \cdot \sum_{t'=\tau}^{\tau+\iota-1} 2G/(\sigma\tau) = 2G^2\iota/\tau \ .
\end{aligned}$$

Summing over $\tau = 1, \ldots, t-\iota$, we get

$$|\mathbb{E}[\mathrm{B}]| \le 2G^2\iota(\log t + 1) \ .$$

Putting together,

$$\begin{aligned}
\mathbb{E}[R_t(w_\Pi^*)] &= \mathbb{E}[\mathrm{A} + \mathrm{B} + \mathrm{C} + \mathrm{D}] \\
&\ge \mathbb{E}[\mathrm{A} + \mathrm{B} + \mathrm{D}] \\
&\ge -\big(|\mathbb{E}[\mathrm{A}]| + |\mathbb{E}[\mathrm{B}]| + |\mathbb{E}[\mathrm{D}]|\big) \\
&\ge -(4\bar{F}\epsilon_t(\iota)t + 2G^2\iota(\log t + 1) + 2\iota\bar{F}) \ ,
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{E}[\|w_{t+1} - w\|^2] &\le \frac{1}{\sigma t}\big(\mathbb{E}[\Delta_t] - \mathbb{E}[R_t(w_\Pi^*)]\big) \\
&\le \frac{\big(6 + \log t\big)G^2}{\sigma^2 t} + \frac{2\big(4\bar{F}\epsilon_t(\iota)t + 2G^2\iota(\log t + 1) + 2\iota\bar{F}\big)}{\sigma t} \ .
\end{aligned} \tag{29}$$

We complete the proof of Theorem 10.

D.5 Proofs for Periodic Case

Assume $\{z_\tau\}_{\tau=1}^t$ are block-wise independent according to the partition $\mathcal{P}$. Given the partition $\mathcal{P}$, define

$$\delta_t^{\mathrm{block}} := \sum_{k=1}^K |I_k| \cdot \left\| \Pi - \frac{1}{|I_k|}\sum_{\tau\in I_k} P^\tau \right\|_{\mathrm{TV}} \ .$$

Note we compute the deviation in a block-wise manner. Note $\delta_t^{\mathrm{block}} = t \cdot \delta^{\mathrm{b}}$ with δ^{b} defined in Theorem 9.

Theorem 11 (Block-wise Independent Data, Restatement of Theorem 9)**.** *It holds*

$$\mathbb{E}\big[\|w_{t+1} - w_\Pi^*\|_2^2\big] = O\left(\frac{\log t}{\sigma^2 t} + \frac{|\mathcal{P}|_\infty^2 \log t}{\sigma t} + \frac{\delta_t^{\mathrm{block}}}{\sigma t}\right) , \tag{30}$$

where $O(\cdot)$ *hides dependence on constants,* G *and* $\bar{F}$.

Generally, compared with δ_t defined in Eq. (23), our new notion of deviation can be much smaller for block-structured data. This is especially true when each block of data, as a whole, forms a good estimate of Π, but each data point in the block deviates from Π by a constant amount. The periodic case in Remark 7 examplifies this.

Remark 5 (Extreme 1: Recover Independent Case)**.** *Setting $|\mathcal{P}|_\infty = 1$ and $\delta_t^{\text{block}} = \delta_t$ in Eq.* (23) *we recover the usual rate under independence assumption (Theorem 7).*

Remark 6 (Extreme 2: Fail to Recover Arbitrary Distribution Case)**.** *If we allow arbitrary dependence in the whole sequence $\gamma = \{z_\tau\}_{\tau=1}^t$, then we can only set $|\mathcal{P}|_\infty = t$ and the bound is useless.*

Remark 7 (The Gain from Block Structure)**.** *Although Theorem 7 applies to block-structure data, we obtain significant improvement in Theorem 11.*

Consider the periodic case where each block is of length q and blocks are i.i.d. At the start of block I_k, we draw a sample from Π, i.e., $z_{t_k} \sim \Pi$, and then let rest of the z_τ's in that block equal z_{t_k}. In this case $\delta_t^{\text{block}} = 0$ because the marginal of every z_τ is exactly Π. Then the bound in Theorem 11 becomes

$$\frac{q^2 \log t}{t} \ , \tag{31}$$

which converges to zero at the rate q^2/t. However, the bound in Theorem 7 fails to converge. To see this let us estimate δ_t in Eq. (23)*. Consider some τ in the interval I_k. If $\tau \neq t_k$, then the conditional distribution $P^\tau(\cdot \,|\, z_{1:\tau-1})$ is a point mass on z_{t_k}. If $\tau = t_k$ then $P^\tau(\cdot \,|\, z_{1:\tau-1}) = \Pi$. Let $c = \sup_z \|\delta_z - \Pi\|_{\text{TV}}$. The quantity c is positive unless Π is a point mass. Then*

$$\delta_t = \sup_{z_1,\ldots,z_t} \sum_{k=1}^K \sum_{\tau \in I_k} \|P^\tau(\cdot \,|\, z_{1:\tau-1}) - \Pi\|_{\text{TV}} = \sup_{z_1,\ldots,z_t} \sum_{k=1}^K (q-1)\|\delta_{z_{t_k}} - \Pi\|_{\text{TV}} = K(q-1)c \ ,$$

and then Theorem 7 becomes

$$\frac{\log t}{t} + c \to c \ .$$

Proof of Theorem 11

We decompose the regret by blocks.

$$R_t(w_\Pi^*) = \sum_{k=1}^K \sum_{\tau \in I_k} \big(F(w_\tau, z_\tau) - \phi_\Pi(w_{\tau_k})\big) + \sum_{\tau=1}^t \big(\phi_\Pi(w_\Pi^*) - F(w_\Pi^*, z_\tau)\big) \tag{32}$$

$$+ \sum_{k=1}^K \sum_{\tau \in I_k} \big(\phi_\Pi(w_{\tau_k}) - \phi_\Pi(w_\Pi^*)\big) \ . \tag{33}$$

Rewrite the first sum by adding and then subtracting the term $\sum_{k=1}^K \sum_{\tau \in I_k} F(w_{\tau_k}, z_\tau)$.

$$\sum_{k=1}^K \sum_{\tau \in I_k} \big(F(w_\tau, z_\tau) - \phi_\Pi(w_{\tau_k})\big)$$

$$= \sum_{k=1}^K \Bigg(\underbrace{\sum_{\tau \in I_k} \big(F(w_\tau, z_\tau) - F(w_{\tau_k}, z_\tau)\big)}_{:=B_k} \Bigg) \tag{I}$$

$$+ \sum_{k=1}^K \Bigg(\underbrace{\sum_{\tau \in I_k} \big(F(w_{\tau_k}, z_\tau) - \phi_\Pi(w_{\tau_k})\big)}_{:=A_k} \Bigg) \ . \tag{II}$$

Bounding A_k. Use a conditioning argument. The key is, conditional on $\mathcal{F}_{\tau_k-1}$, the iterate w_{τ_k} is deterministic and the distribution of z_τ is P^τ due to block-wise independence.

$$
\begin{aligned}
|\mathbb{E}[A_k]| &= \left|\sum_{\tau\in I_k}\mathbb{E}[\mathbb{E}[F(w_{\tau_k},z_\tau)-\phi_\Pi(w_{\tau_k})|\mathcal{F}_{\tau_k-1}]]\right| \\
&\le \mathbb{E}\left[\left|\sum_{\tau\in I_k}\mathbb{E}[F(w_{\tau_k},z_\tau)-\phi_\Pi(w_{\tau_k})|\mathcal{F}_{\tau_k-1}]\right|\right] \\
&= \mathbb{E}\left[\left|\sum_{\tau\in I_k}\int_{\mathcal{Z}}F(w_{\tau_k},z)P^\tau(\mathrm{d}z)-\int_{\mathcal{Z}}F(w_{\tau_k},z)\Pi(\mathrm{d}z)\right|\right] \\
&\le 2\bar{F}\cdot\left\|\sum_{\tau\in I_k}\left(P^\tau-\Pi\right)\right\|_{\mathrm{TV}} .
\end{aligned}
$$

Then sum over $k=1,\dots,K$, and we have

$$
\left|\mathbb{E}[\mathrm{I}]\right| \le \sum_{k=1}^{K}\left|\mathbb{E}[A_k]\right| \le 2\bar{F}\cdot\sum_{k=1}^{K}\left\|\sum_{\tau\in I_k}(P^\tau-\Pi)\right\|_{\mathrm{TV}} \le 2\bar{F}\delta_t^{\mathrm{block}} .
$$

Bounding B_k. Using Lemma 6 and Eq. (28), we have

$$
\begin{aligned}
|\mathbb{E}[B_k]| &\le G\cdot\sum_{\tau\in I_k}\mathbb{E}\left[\|w_\tau-w_{\tau_k}\|\right] \\
&\le G\cdot\sum_{\tau\in I_k}G(\tau-\tau_k)/\tau_k \\
&\le G\cdot\sum_{\tau\in I_k}G(\tau_{k+1}-\tau_k)/\tau_k \\
&= G^2(\tau_{k+1}-\tau_k)^2/\tau_k \le G^2|\mathcal{P}|_\infty^2/\tau_k .
\end{aligned}
$$

Then sum over $k=1,\dots,K$, and we have

$$
\left|\mathbb{E}[\mathrm{II}]\right| \le \sum_{k=1}^{K}\left|\mathbb{E}[B_k]\right| = G^2|\mathcal{P}|_\infty^2\cdot\sum_{k=1}^{K}\frac{1}{\tau_k} \le G^2|\mathcal{P}|_\infty^2\cdot\sum_{\tau=1}^{t}\frac{1}{\tau} \le G^2|\mathcal{P}|_\infty^2(\log t+1) .
$$

It can be shown the second sum in the regret decomposition (Eq. (32)) is upper bounded by $2\bar{F}\delta_t^{\mathrm{block}}$. The third sum is ≥ 0. Using Lemma 1 we get

$$
\mathbb{E}[\|w_{t+1}-w\|^2] \le \frac{1}{\sigma t}\left(\mathbb{E}[\Delta_t]-\mathbb{E}[R_t(w_\Pi^*)]\right) \le \frac{(6+\log t)G^2}{\sigma^2 t}+\frac{2\left(4\bar{F}\delta_t^{\mathrm{block}}+G^2|\mathcal{P}|_\infty^2(\log t+1)\right)}{\sigma t} . \tag{34}
$$

We complete the proof of Theorem 11.

E Proofs for PACE

E.1 Proof of Theorem 5 and Theorem 6

We show convergence of β under different input assumption. Recall β^t is the pacing multiplier generated by PACE, and β^* is the solution to the optimization problem Eq. (4). The vector γ is the sequence of items.

Theorem 12 (Restatement of Theorem 5 and Theorem 6)**.** *For the independent case, i.e., $\gamma\sim Q$ and $Q\in\mathsf{C}^{\mathrm{ID}}(\delta)$, it holds for $t\ge 1$*

$$
\mathbb{E}[\|\beta^{t+1}-\beta^*\|^2] \le \frac{(6+\log t)G^2}{\sigma^2 t}+\frac{8\bar{F}}{\sigma}\delta . \tag{35}
$$

and for $t \geq 3$,

$$\frac{1}{t}\sum_{\tau=1}^{t}\mathbb{E}[\|\beta^\tau - \beta^*\|^2] \leq \frac{G^2}{\sigma^2}\left(6(1+\log t) + \frac{(\log t)^2}{2}\right) + (8\bar{F}/\sigma)\cdot\delta\ . \tag{36}$$

For the erogdic case, i.e., $\gamma \sim Q$ and $Q \in \mathsf{C}^{\mathrm{E}}(\delta, \iota)$, it holds for $t \geq 1$,

$$\mathbb{E}[\|\beta^{t+1} - \beta^*\|^2] \leq \frac{\big(6+\log t\big)G^2}{\sigma^2 t} + \frac{2\big(4\bar{F}\delta t + 2G^2\iota(\log t + 1) + 2\iota\bar{F}\big)}{\sigma t} = \tilde{O}\left(\delta + \frac{\iota}{t}\right). \tag{37}$$

and for $t \geq 3$,

$$\frac{1}{t}\sum_{\tau=1}^{t}\mathbb{E}[\|\beta^\tau - \beta^*\|^2] = \tilde{O}\left(\delta + \frac{\iota}{t}\right)\ . \tag{38}$$

For the periodic case, i.e., $\gamma \sim Q$ and $Q \in \mathsf{C}^{\mathrm{P}}(q)$, it holds for $t \geq 1$

$$\mathbb{E}[\|\beta^{t+1} - \beta^*\|^2] \leq \frac{(6+\log t)G^2}{\sigma^2 t} + \frac{2G^2q^2(\log t + 1)}{\sigma t} = \tilde{O}(q^2/t)\ . \tag{39}$$

and for $t \geq 3$,

$$\frac{1}{t}\sum_{\tau=1}^{t}\mathbb{E}[\|\beta^\tau - \beta^*\|^2] = \tilde{O}(q^2/t)\ . \tag{40}$$

Proof of Theorem 12. Set $\sigma = 1/n$ and $G = \bar{F} = |v|_\infty$. Eq. (35) follows by Theorem 7 and specifically Eq. (25). The next inequality Eq. (36) follows by (Xiao, 2010, Corollary 4): for $t \geq 3$,

$$\frac{1}{t}\sum_{\tau=1}^{t}\frac{(6+\log\tau)G^2}{\tau\sigma^2} \leq \frac{1}{t}\left(6(1+\log t) + \frac{(\log t)^2}{2}\right)\frac{G^2}{\sigma^2}\ .$$

Eq. (37) follows by Theorem 10 and specifically Eq. (29). Following (Xiao, 2010, Corollary 4), we have for $t \geq 3$,

$$\begin{aligned}
&\sum_{\tau=1}^{t}\mathbb{E}[\|\beta^\tau - \beta^*\|^2] \\
&\leq \frac{G^2}{\sigma^2}\left(6(1+\log t) + \frac{(\log t)^2}{2}\right) + 8\bar{F}\cdot\delta t + \frac{4G^2}{\sigma}\left(6(1+\log t) + \frac{(\log t)^2}{2}\right) + \frac{4\bar{F}}{\sigma}(1+\log t)\cdot\iota \\
&= \tilde{O}(\delta t + \iota)\ ,
\end{aligned}$$

and thus Eq. (38) holds.

The inequality Eq. (39) follows from Theorem 11 and specifically Eq. (34). For the inequality Eq. (40), apply the same strategy: for $t \geq 3$,

$$\begin{aligned}
&\sum_{\tau=1}^{t}\mathbb{E}[\|\beta^\tau - \beta^*\|^2] \\
&\leq \frac{G^2}{\sigma^2}\left(6(1+\log t) + \frac{(\log t)^2}{2}\right) + \frac{2G^2}{\sigma^2}\left(6(1+\log t) + \frac{(\log t)^2}{2}\right)\cdot q^2 = \tilde{O}(q^2)\ .
\end{aligned}$$

□

E.2 Proof of Lemma 3 and Lemma 4

Proof of Lemma 3

Lemma 3 follows from (Gao et al., 2021, Theorem 3 and 4).

Proof of Lemma 4

Define the hindsight *average* equilibrium utility $u_i^\gamma := (1/t) \cdot U_i^\gamma$. Although results in Gao et al. (2021) were stated for i.i.d. case, the proof in fact goes through for nonstationary input distributions.

Bounding $\mathbb{E}[\|\bar{u}^t - u^\gamma\|^2]$. For the first inequality we use the proof of (Gao et al., 2021, Theorem 6). Follow that paper, we define $r_i^t = \max\{0, \bar{u}_i^t - u_i^\gamma\}$. In Theorem 6 the authors show

$$\mathbb{E}[(r_i^t)^2] \leq C_{r,1} \cdot \mathbb{E}[\|\beta^{t+1} - \beta^*\|^2] + C_{r,2} \cdot \left(\frac{1}{t}\sum_{\tau=1}^{t}\mathbb{E}[\|\beta^\tau - \beta^*\|^2]\right).$$

In particular, the constant $C_{r,1}$ comes from the constant C in (Gao et al., 2021, Theorem 4) and $C_{r,2}$ comes from (Gao et al., 2021, Equation (11))

Bounding $\mathrm{Reg}_{i,t}$. Note $\mathrm{Reg}_{i,t} = t \cdot (u_i^\gamma - \bar{u}_i^t) \leq t \cdot r_i^t$. Then we use Cauchy-Schwarz.

$$\mathbb{E}\big[\mathrm{Reg}_{i,t}\big] \leq t\mathbb{E}[r_i^t] \leq t\sqrt{\mathbb{E}[(r_i^t)^2]}\,.$$

Bounding $\mathrm{Envy}_{i,t}$. For the second inequality we use the proof of Theorem 6 in the same paper. Following that paper, we define

$$\rho_i^t = (n/t) \cdot \max_{k\in[n]}\Big\{\langle v_i(\gamma), x_k\rangle - \langle v_i(\gamma), x_i\rangle\Big\}\,.$$

During the course of proving Theorem 6, the authors show

$$\mathbb{E}[(\rho_i^t)^2] \leq n^2\left(C_{e,1} \cdot \mathbb{E}[\|\beta^{t+1} - \beta^*\|^2] + C_{e,2} \cdot \Big(\frac{1}{t}\sum_{\tau=1}^{t}\mathbb{E}[\|\beta^\tau - \beta^*\|^2]\Big)\right).$$

In particular, the constant $C_{e,1}$ comes from (Gao et al., 2021, Theorem 4, Equations (5) and (13)) and $C_{e,2}$ comes from (Gao et al., 2021, Equations (5), (13) and (15))

Then using Cauchy-Schwarz,

$$\mathbb{E}\big[\mathrm{Envy}_{i,t}(\gamma)\big] = \mathbb{E}\big[(t/n) \cdot \rho_i^t\big] \leq (t/n)\sqrt{\mathbb{E}\big[(\rho_i^t)^2\big]}\,.$$

This completes the proof of Lemma 4.

F Experiments

We conduct experiments on a market (a matrix of buyers' valuations on items) generated from the MovieLens dataset (Harper and Konstan, 2016) with $n = 100$ buyers and $m = 300$ items. The process of turning the MovieLens dataset into the market instance is described in Kroer et al. (2021). Here, we briefly describe the experiment settings. For more details on the experiment settings as well as all code and data to replicate the results, please refer to the Supplementary Material.

We generate item arrivals from the following data input models:

- i.i.d.: Every item $\theta^t \in [m]$ is sampled independently from a fixed distribution $s^0 \in \Delta_m$ (an m-dimensional probability vector).
- Mild corruption: $\theta^t \sim s^t$, where $s^t \in \Delta_m$ is a distribution such that $\|s^t - s^0\|_1 = \Theta(1/t)$ for all t. Here, s^t is generated by randomly perturbing each coordinate of s^0 followed by normalization.
- Markov: $(\theta^t)_{t\geq 1}$ is sampled from an irreducible Markov chain starting from an initial distribution s^0. It is a special case of ergodic input. Here, the Markov chain is given by a $m \times m$ transition matrix (each row sums to 1), which we generated randomly (and row-wise normalized). In this case, the "reference" item arrival distribution is the stationary distribution of this Markov chain which is in general different from the initial distribution.
- Periodic: The period length is $\ell = 100$. Let $(s^k)_{k\in[\ell]}$ be a set of distributions (probability vectors). Here, each s^k is sampled randomly and normalized. The item arrivals of each period is generated by sampling from each s^k followed by a random permutation over the ℓ sampled items.

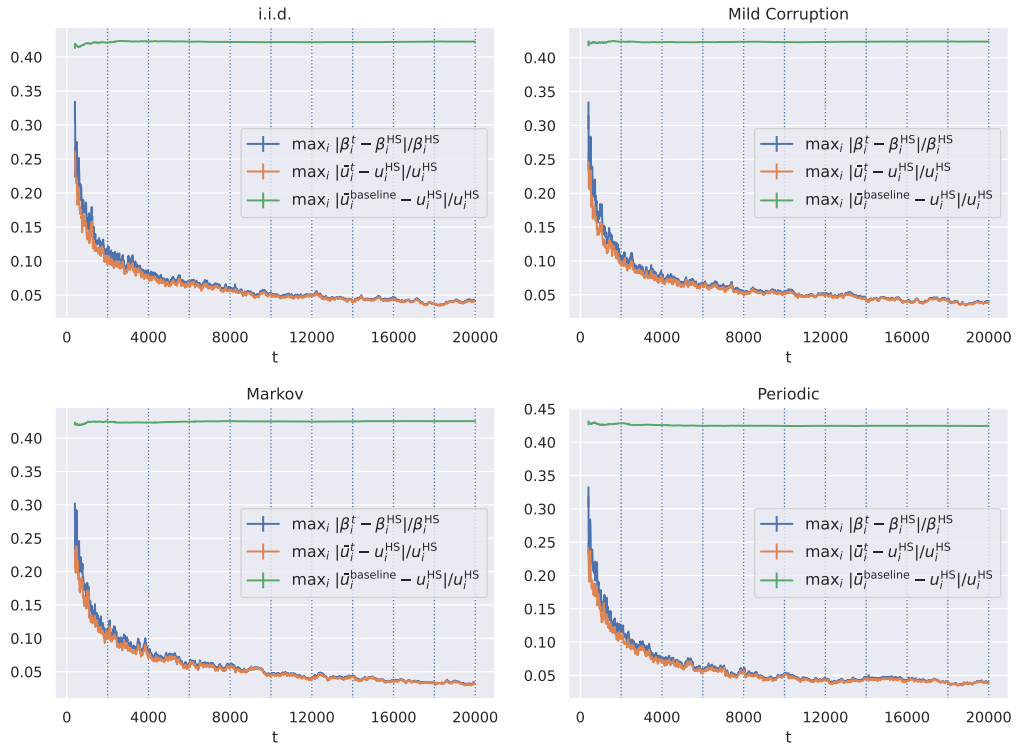

Figure 1: Performance of PACE for item arrivals under different data input models. All error measures are averaged across 10 repeated experiments. The mean and standard errors of the error measures are plotted, where the standard error bars are too small and hence invisible. Here, $\bar{u}_i^{\text{baseline}}$ are the buyers' time-averaged utilities under a "proportional-share" baseline solution.

For each (fixed) data input model, we generate 10 sample paths of item arrivals and run PACE for $T = 200n = 20000$ time steps on each sample path. Then, we measure the convergence of the pacing multipliers and time-averaged cumulative utilities to their hindsight equilibrium values. More specifically, we record the following relative differences: $\max_i\{|\beta_i^t - \beta_i^{\text{HS}}|/\beta_i^{\text{HS}}\}$ and $\max_i\{|\bar{u}_i^t - u_i^{\text{HS}}|/u_i^{\text{HS}}\}$, where HS denote the hindsight equilibrium values of the "sample-path" market determined by the realized item arrivals. Equivalently, u^{HS} and β^{HS} are optimal solutions of the hindsight convex programs (1) and (2), respectively. We also measure the performance of a proportional-share baseline solution that divides each arriving item among all buyers proportionally w.r.t. their budgets: for an arrived item θ^t, each buyer i gets B_i amount of it and receives utility $B_i v_i(\theta^t)$ (in this paper, the buyers' budgets are $B_i = 1/n$ for all i). We compute the means and standard errors of the error measures across the 10 sample paths and plot them in Figure 1.

As can be seen, for all data input models, the pacing multipliers and buyers' time-averaged utilities converge to their respective hindsight values and quickly outperform the baseline proportional-share solution. Similar convergence behavior can also be observed when the error metrics are w.r.t. to the true equilibrium values β^*, u^* instead of the hindsight values.

We further conduct the following experiment to demonstrate the effect of nonstationarity on pacing multiplier convergence. More specifically, we generate 10 sample paths from each of the following item arrival settings: i.i.d. distributions ($\theta_t \sim_{\text{i.i.d.}} s$), perturbed distributions (small δ), perturbed distributions (large δ), where δ is the overall difference between the sequences of i.i.d. distributions and perturbed distributions, as in (**??**). We then run PACE on each sample path to obtain β_i^t and $\bar{u}_i^t$ values. For each setting and each time step, we plot mean values and standard error bars of the relative error metrics, where β_i^* and u_i^* are the equilibrium pacing multipliers (utility prices) and utilities of the market with supplies being the distribution s. As can be seen, for perturbed distribution settings, PACE is able to bring β^t and $\bar{u}_i^t$ close to their equilibrium values, while the

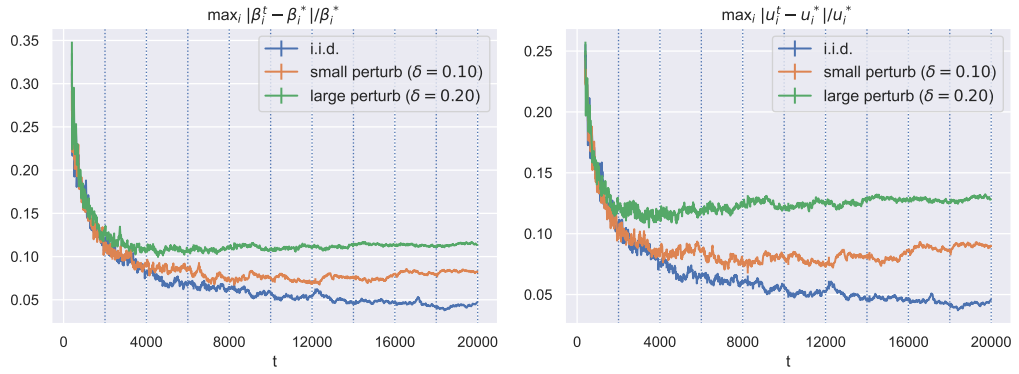

Figure 2: Convergence of pacing multipliers β^t and cumulative utilities $\bar{u}_i^t$ under i.i.d. and perturbed distributions

convergence degrade as the perturbation amount δ increases. Recall that, in terms of cumulative utility, the proportional-share baseline solution gives a relative error of around 0.45.