

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 [Birnbbaum 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 [Birnbbaum 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. [Batoni 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 man-

agement strategies for Internet advertising auctions, with strong revenue and incentive guarantees (see, e.g., [Conitzer et al \(2019, 2021\)](#); [Balseiro and Guri \(2019\)](#)). 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 [Cao 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 $F = (n, m, B, v)$. The market consists of n buyers and m items. We assume each buyer has a budget of B_i . We use $\{1, \dots, m\}$ to represent the set of items, each of unit supply. The matrix $v = (v_1, \dots, v_n) \in (\mathbb{R}_+^m)^n$ consists of valuations, with v_i^j being the valuation of item j from buyer i . For buyer i , an allocation of items, $x_i \in \mathbb{R}_+^m$, gives a utility of $u_i(x_i) := \langle v_i, x_i \rangle := \sum_{j=1}^m v_i^j x_i^j$. 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:*

$$D_i(p) := \arg \max \{ \langle v_i, x_i \rangle : x_i \geq 0, \langle p, x_i \rangle \leq B_i \}. \quad (17)$$

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

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

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 $\sum_{i=1}^n x_i \leq \mathbf{1}_m$. A celebrated result is the Eisenberg-Gale convex program, which provides an equivalent characterization of ME.

$$\max_{x_1, \dots, x_n} \sum_{i=1}^n B_i \log \langle v_i, x_i \rangle \quad \text{s.t.} \quad \sum_{i=1}^n x_i^j \leq 1 \quad \forall j \in [m], \quad x_i \in \mathbb{R}_+^m \quad \forall i \in [n]. \quad (18)$$

That is, we maximize the sum of logarithmic utilities under the supply constraint. The solution to the primal problem $x^* = (x_1^*, \dots, x_n^*)$ along with the vector of dual variables p^* yields a market equilibrium.

The hindsight allocation Eq. (III) is just the EG program of the linear Fisher market $F_A = (n, t, B, v)$ where entries of the valuation matrix v are defined by $v_i^j = v_i(\theta^j)$ 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 $d\bar{Q}/d\mu$ 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 [Kia0 \(2010\)](#), 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 [Kia0 \(2010, Algorithm 1\)](#) 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 (III). 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 $L : \mathbb{R}_+^{n \times t} \times \mathbb{R}_+^n \times \mathbb{R}_+^n \times \mathbb{R}_+^t \rightarrow \mathbb{R}$ is given by

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

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

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

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:

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

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 (III):

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

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

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

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 (21) 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 (21). 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 (21). Define the corresponding optimal p^* in (2) 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

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

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

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

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 4) are exactly DA(G, Ψ, γ) (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 $[(1+\delta_n)^{-1} n^{-1} \mathbf{1}_n, (1+\delta_0) \mathbf{1}_n]$ 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 $G(\beta, \theta)$ is exactly the time- τ realized utility of agent i in PACE, that is, $g_{\tau,i} = v_i(\theta^\tau) \mathbf{1}\{i = i^\tau\} = u_i^\tau$. 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:

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

(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 $\bar{g}_{\tau,i} = \bar{u}_i^\tau$:

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

□

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 (21).

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

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

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. \square

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 ergodic case, i.e., $\gamma \sim Q$ and $Q \in \mathcal{C}^{\text{E}}(\delta, \iota)$, 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 $C_{E,1} = n^2|v|_\infty^2(6 + \log t)$, $C_{E,2} = 4n(|v|_\infty^2(1 + \log t) + |v|_\infty)$ and $C_{E,3} = 8n|v|_\infty$.

For the periodic case, i.e., $\gamma \sim Q$ and $Q \in \mathcal{C}^{\text{P}}(q)$, 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 4. For the second inequality, set additionally $\delta^b = 0$ and $|\mathcal{P}|_\infty = q$ in Theorem 4. \square

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 $t \geq 1$,*

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

and

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

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

Finally, we relate results in Appendix C3 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 $t \geq 1$,*

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

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

and

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

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

D Proofs for Nonstationary Dual Averaging

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

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

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:

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

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

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

It follows that all lower bounds for the regret $\mathbb{E}[R_t(w_\Pi^*)]$ 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 $\|w_{t+1} - w_\Pi^*\|^2$ 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 $P^{\tau'}(\cdot | z_{1:0}) = P^{\tau'}$, the marginal distribution of $z_{\tau'}$.

Define the progressive deviation from Π

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

If the data are independent, then it holds

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

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*

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

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 4, 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:

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

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

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

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

Proof of Lemma 5. Write

$$R_t(w_\Pi^*) = \sum_{\tau=1}^t (F(w_\tau, z_\tau) - \phi_\Pi(w_\tau)) \quad (\text{I})$$

$$+ \sum_{\tau=1}^t (\phi_\Pi(w_\Pi^*) - F(w_\Pi^*, z_\tau)) \quad (\text{II})$$

$$+ \sum_{\tau=1}^t (\phi_\Pi(w_\tau) - \phi_\Pi(w_\Pi^*)). \quad (\text{III})$$

By optimality of w_Π^* we have III ≥ 0 .

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

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 $P^\tau(\cdot | z_{1:\tau-1})$ by definition.

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

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

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

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

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

Combining we get

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

This completes the proof of Lemma [5](#). □

D.3 Theorem Statements for Markov Case and Periodic Case

Results for other input types, $C^E(\delta, \iota)$ and $C^P(q)$, 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 $1 \leq \iota \leq t-1$:*

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

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

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

Moreover, the rate $\tilde{O}(\epsilon_t(\iota) + \iota/t)$ applies to $\mathbb{E}[\|w_{t+1} - w_\gamma^*\|_2^2]$ and $\mathbb{E}[\|w_\gamma^* - w_\Pi^*\|_2^2]$ by the deduction in Appendix [D.7](#).

Remark 4 (Comparison with EMD [Duchi et al \(2012\)](#)). Now we specialize Theorem [5](#) to the setting of Remark [4](#), 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 $f^* = \min \{f(w) = \mathbb{E}_\Pi[F(w; \xi)] \mid w \in \mathcal{W}\}$ 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 $\epsilon_t(\iota) \leq M\rho^t$ for some $M > 0$ and $\rho \in [0, 1)$, then the EMD algorithm produces iterates that satisfy the following convergence rate [8](#)

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

As in Remark [4](#), 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

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

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 $K \geq 1$. Let $\{1 = \tau_1 < \tau_2 < \dots < \tau_{K+1} = t\}$ 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 $|I_k| := \tau_{k+1} - \tau_k \geq 1$ 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*

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

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

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

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

Let us briefly we comment on the dependence on $|\mathcal{P}|_\infty$. 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 $\delta^b = 0$ 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.

Define the ι -step deviation from stationarity

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

³In 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) .

An equivalent quantity is the ϵ -mixing time [Duchi et al \(2012\)](#)

$$\iota_{\text{mix}}(\epsilon) := \min \left\{ \iota : 1 \leq \iota \leq t-1, \sup_{z_1, \dots, z_\iota} \sup_{\tau=1, \dots, t-\iota} \|P^{\tau+\iota}(\cdot | z_{1:\tau}) - \Pi\|_{\text{TV}} \leq \epsilon \right\}.$$

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 $t \geq 1$ and any $1 \leq \iota \leq t-1$,*

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

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.

$$R_t(w_{\Pi}^*) = \sum_{\tau=1}^{t-\iota} \left((F(w_\tau, z_{\tau+\iota}) - F(w_{\Pi}^*, z_{\tau+\iota})) - (\phi_{\Pi}(w_\tau) - \phi_{\Pi}(w_{\Pi}^*)) \right) \quad (\text{A})$$

$$+ \sum_{\tau=1}^{t-\iota} (F(w_{\tau+\iota}, z_{\tau+\iota}) - F(w_\tau, z_{\tau+\iota})) \quad (\text{B})$$

$$+ \sum_{\tau=1}^{t-\iota} (\phi_{\Pi}(w_\tau) - \phi_{\Pi}(w_{\Pi}^*)) \quad (\text{C})$$

$$+ \sum_{\tau=1}^{\iota} (F(w_\tau, z_\tau) - F(w_{\Pi}^*, z_\tau)). \quad (\text{D})$$

By optimality of w_{Π}^* we have C ≥ 0 . By boundedness of F we get |D| $\leq 2\iota\bar{F}$. Remains to handle A and B. We will show

$$" A \leq \epsilon_t(\iota)t, \quad B \leq \iota " .$$

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

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

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

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

Bounding B. 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*

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

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

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

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

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

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

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

Putting together,

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

and

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

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

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

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

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

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

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 exemplifies 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 4).

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 4 applies to block-structure data, we obtain significant improvement in Theorem 5.

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 5 becomes

$$\frac{q^2 \log t}{t}, \quad (31)$$

which converges to zero at the rate q^2/t . However, the bound in Theorem 4 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

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

and then Theorem 4 becomes

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

Proof of Theorem 5

We decompose the regret by blocks.

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

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

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

$$\begin{aligned} & \sum_{k=1}^K \sum_{\tau \in I_k} (F(w_\tau, z_\tau) - \phi_\Pi(w_{\tau_k})) \\ &= \sum_{k=1}^K \left(\underbrace{\sum_{\tau \in I_k} (F(w_\tau, z_\tau) - F(w_{\tau_k}, z_\tau))}_{:=B_k} \right) \end{aligned} \quad (\text{I})$$

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

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

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

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

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

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

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

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

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

It can be shown the second sum in the regret decomposition (Eq. (32)) is upper bounded by $2\bar{F} \delta_t^{\text{block}}$. The third sum is ≥ 0 . Using Lemma 11 we get

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

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 $Q \in C^{\text{ID}}(\delta)$, it holds for $t \geq 1$*

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

and for $t \geq 3$,

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

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

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

and for $t \geq 3$,

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

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

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

and for $t \geq 3$,

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

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

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

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

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

During the course of proving Theorem 6, the authors show

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

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,

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

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.

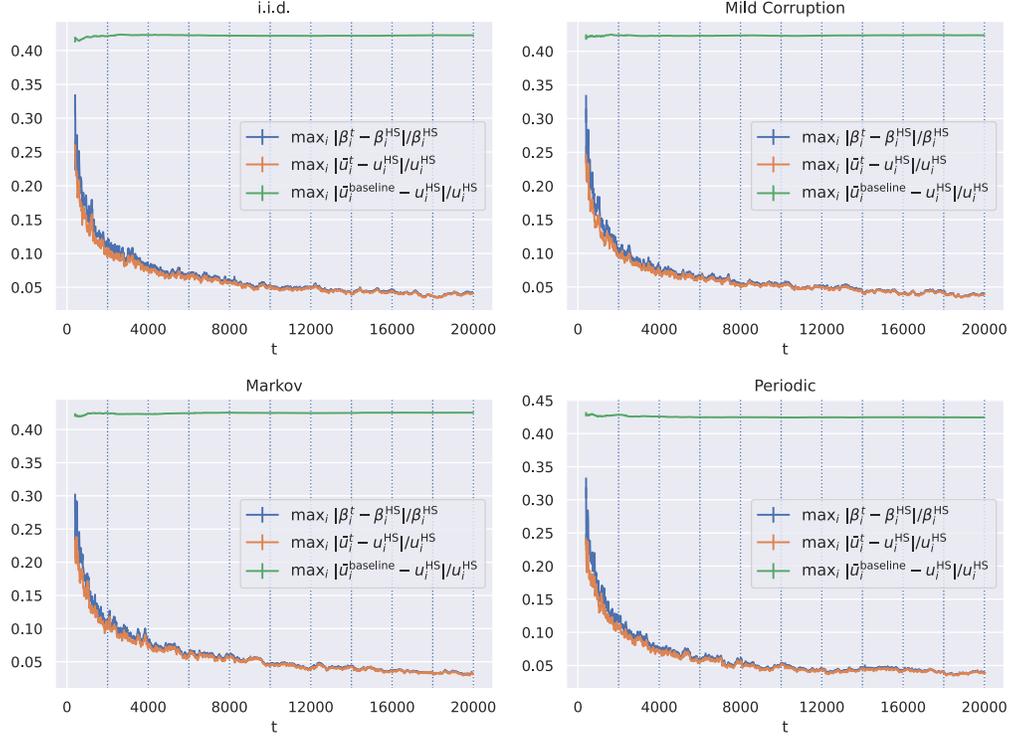


Figure 1: Performance of PACE for item arrivals under different data input models. All error measures are averaged across 10 repeated experiments. The mean and standard errors of the error measures are plotted, where the standard error bars are too small and hence invisible. Here, $\bar{u}_i^{\text{baseline}}$ 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 (III) and (II), 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

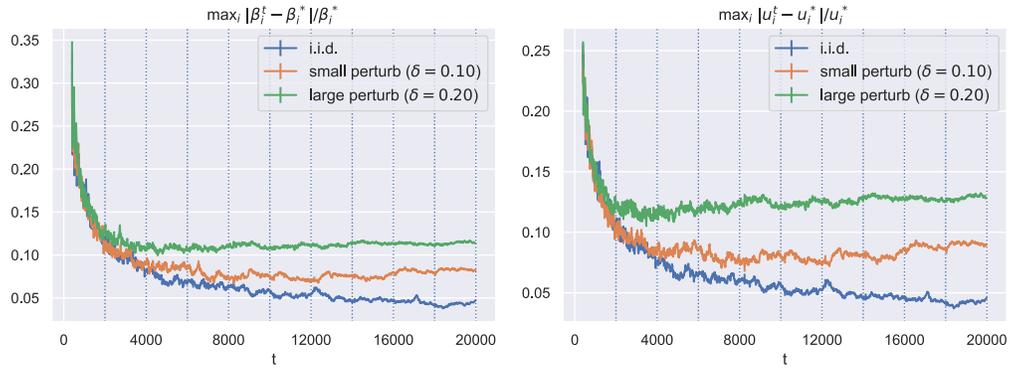


Figure 2: Convergence of pacing multipliers β^t and cumulative utilities \bar{u}_i^t under i.i.d. and perturbed distributions

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.