
Online Market Equilibrium with Application to Fair Division

Yuan Gao
Columbia University
gao.yuan@columbia.edu

Alex Peysakhovich
Facebook AI Research
alex.peys@gmail.com

Christian Kroer
Columbia University
christian.kroer@columbia.edu

Abstract

Computing market equilibria is a problem of both theoretical and applied interest. Much research to date focuses on the case of static Fisher markets with full information on buyers' utility functions and item supplies. Motivated by real-world markets, we consider an online setting: individuals have linear, additive utility functions; items arrive sequentially and must be allocated and priced irrevocably. We define the notion of an online market equilibrium in such a market as time-indexed allocations and prices which guarantee buyer optimality and market clearance in hindsight. We propose a simple, scalable and interpretable allocation and pricing dynamics termed as PACE. When items are drawn i.i.d. from an unknown distribution (with a possibly continuous support), we show that PACE leads to an online market equilibrium asymptotically. In particular, PACE ensures that buyers' time-averaged utilities converge to the equilibrium utilities w.r.t. a static market with item supplies being the unknown distribution and that buyers' time-averaged expenditures converge to their per-period budget. Hence, many desirable properties of market equilibrium-based fair division such as envy-freeness, Pareto optimality, and the proportional-share guarantee are also attained asymptotically in the online setting. Next, we extend the dynamics to handle quasilinear buyer utilities, which gives the first online algorithm for computing first-price pacing equilibria. Finally, numerical experiments on real and synthetic datasets show that the dynamics converges quickly under various metrics.

1 Introduction

A market is said to be in equilibrium when supply is equal to demand. Computing prices and allocations which constitute a market equilibrium (ME) has long been a topic of interest [17, 20, 28, 31, 38, 43]. Most existing work focuses on the case of static markets. However, in this paper we consider the case of online markets where items arrive sequentially. We consider the extension of market equilibrium to this setting and provide market dynamics which quickly converge to an equilibrium in the case of online Fisher markets.

In static Fisher markets there is a fixed supply of each item, individual preferences are linear, additive, and items are divisible (or equivalently, randomization is allowed so individuals can purchase not just items but lotteries over items). In general, finding market equilibria is a hard problem [14, 39, 47]. However, in static linear Fisher markets, equilibrium prices and allocations can be computed via solving the Eisenberg-Gale (EG) convex program [22, 37].

We consider an online extension of Fisher markets where buyers are constantly present but items arrive one-at-a-time. Buyers' budgets are per-period and represent their respective 'bidding powers' instead of being binding constraints. We extend the definition of market equilibrium to the online setting: online equilibrium allocations and prices are time-indexed and, when averaged across time, form an equilibrium in a corresponding static Fisher market where item supplies are proportional

to item arrival probabilities. Due to the stochastic nature of online Fisher markets, any online algorithm can only attain an online market equilibrium *asymptotically*, that is, the allocations and prices approximately satisfy the equilibrium conditions after running the algorithm for a long time.

We propose market dynamics that find these equilibria in an online fashion based on the *dual averaging* algorithm applied to a reformulation of the dual of the EG convex program. We refer to this mechanism as **PACE** (Pace According to Current Estimated utility). In PACE, each buyer is assigned a utility *pacing* multiplier at time 0. When an item arrives, the individual with the highest adjusted utility (its valuation times the multiplier) receives that item and pays a price equal to its adjusted utility. The pacing multipliers of all individuals are then adjusted according to a closed-form rule which is given by the time average of the subgradient of the dual of the EG program. Intuitively, the pacing multipliers of those that did not receive the item go up while the receiver’s typically (but not always) goes down. We show that PACE yields item allocations and prices that satisfy various equilibrium properties asymptotically, for example no-regret and envy-freeness.

One important application of market equilibrium is fair allocation using the *competitive equilibrium from equal incomes* (CEEI) mechanism [12, 46]. In CEEI, each individual is given an endowment of faux currency and reports her valuations for items; then, a market equilibrium is computed and the items are allocated accordingly. However, many fair division problems are online rather than static. These include the allocation of impressions to content in certain recommender systems [34], workers to shifts, donations to food banks [2], scarce compute time to requestors [25, 29, 40], or blood donations to blood banks [32]. Similarly, online advertising can also be thought of as the allocation of impressions to advertisers via a market though with a budget of real money rather than faux currency. In the static CEEI case with linear additive preferences, the resulting equilibrium outcomes (i.e. results of the EG program) have been described as “perfect justice” [3]. In the online case, PACE achieves the same fair allocations as CEEI asymptotically. See Appendix A for more related work in the areas of (static and online) equilibrium computation and fair division.

We evaluate PACE experimentally in several market datasets. Convergence to good outcomes happens quickly in experiments. Taken together our results, we conclude that PACE is an attractive algorithm for both computing online market equilibria and online fair division.

Main contributions. We consider the problem of allocating and pricing sequentially arriving items to n buyers. This setting is termed as an *online Fisher market*. Given a sequence of item arrivals, we define an online market equilibrium as the items’ allocations and prices that, in hindsight, ensure buyer optimality and market clearance. We propose the PACE dynamics, which can be viewed as a nontrivial instantiation of the dual averaging algorithm on a reformulation of the dual of the Eisenberg-Gale convex program. Leveraging the convergence theory of dual averaging, we show that, when item arrivals are drawn from an (unknown) underlying distribution s , possibly over an infinite/continuous item space, PACE ensures the following.

- The pacing multipliers generated by PACE converge to the static equilibrium *utility prices*. Here, “static” means w.r.t. to an underlying static Fisher market.
- Buyers’ time-averaged utilities converge to the static equilibrium utilities.
- Buyers’ time-averaged expenditures converge to their respective budgets.

These convergences are all in mean square with rates $O((\log t)/t)$, $O((\log t)/t)$ and $O((\log t)^2/t)$, respectively, where the constants in these rates involve moderate polynomials of n . In this way, PACE generates allocations and prices that constitute an online market equilibrium in the limit. In particular, the allocations and prices ensure that the allocation is Pareto optimal, and buyers have no regret, no envy, and get at least their proportional share asymptotically. We also extend PACE to the case of quasilinear buyer utilities, which yields the first online algorithm for computing first-price pacing equilibria. Finally, numerical experiments suggest that PACE converges much faster than its theoretical rates in terms of pacing multipliers, utilities and expenditures.

2 Static and Online Fisher Markets

Static Fisher markets and equilibria. We first introduce static Fisher markets and their equilibria. Following the recent work [24, §2], we consider a measurable (possibly continuous) item space.

Below are the technical preliminaries for the subsequent online setting. They can be skimmed through and referred back to as needed.

From now on, we define $[k] := \{1, \dots, k\}$ for any $k \in \mathbb{N} := \{0, 1, 2, \dots\}$ and \mathbb{R}_+ (\mathbb{R}_{++} , resp.) as the set of nonnegative (positive, resp.) real numbers. Let $\mathbb{I}\{A\} \in \{0, 1\}$ denote the indicator function of an event A .

- (a) There are n buyers (individuals), each having a budget $B_i > 0$.
- (b) The item space is a *finite* measure space $(\Theta, \mathcal{M}, \mu)$ with $0 < \mu(\Theta) < \infty$. From now on, L^p (and L^p_+ , resp.) denote the set of (nonnegative, resp.) L^p functions on Θ for any $p \in [1, \infty]$ (including $p = \infty$). Below are some concrete special cases for illustration.
 - (i) Finite: $\Theta = [m]$, $\mathcal{M} = 2^{[m]} = \{A : A \subseteq [m]\}$ and $\mu(A) = \sum_{a \in A} \mu(a)$ (all 2^m subsets are measurable and the measure is given by a point mass on each item).
 - (ii) Lebesgue-measurable: μ is the Lebesgue measure on \mathbb{R}^d , \mathcal{M} is the Lebesgue σ -algebra and Θ is a (Lebesgue-)measurable subset of \mathbb{R}^d with positive finite measure. For example, Θ can be a compact subset of \mathbb{R}^d with a nonempty interior.
 - (iii) Countably infinite: $\Theta = \mathbb{N}$ and $\mu(A) = \sum_{a \in A} \mu(a)$ for any $A \subseteq \mathbb{N}$, where $\mu(\mathbb{N}) < \infty$. For example, $\mu(a)$ can be the probability mass of a Poisson distribution, in which case $(\mathbb{N}, \mathcal{M}, \mu)$ is a probability space.
- (c) The *supplies* of items is $s \in L^p_+$, i.e., item $\theta \in \Theta$ has supply $s(\theta)$. Since Θ is compact, it is measurable with a finite measure. For the finite case $\Theta = [m]$, we have $s = (s_1, \dots, s_m) \in \mathbb{R}^m_+$.
- (d) The *valuation* of each buyer i on all items is $v_i \in L^1_+$, i.e., buyer i has valuation $v_i(\theta)$ on item $\theta \in \Theta$. For the finite case $\Theta = [m]$, we have $v_i = (v_{i1}, \dots, v_{im}) \in \mathbb{R}^m_+$.
- (e) For buyer i , an *allocation* of items $x_i \in L^\infty_+$ gives a utility of

$$u_i(x_i) := \langle v_i, x_i \rangle := \int_{\Theta} v_i(\theta) x_i(\theta) d\theta,$$

where the angle brackets are based on the notation of applying a bounded linear functional x_i to a vector v_i in the Banach space L^1 and the integral is the usual Lebesgue integral. For the finite case $\Theta = [m]$, we have $x_i = (x_{i1}, \dots, x_{im}) \in \mathbb{R}^m_+$ and the utility is

$$u_i(x_i) = \langle v_i, x_i \rangle = \sum_j v_{ij} x_{ij},$$

the usual Euclidean vector inner product. We will use $x \in (L^\infty_+)^n$ to denote the aggregate allocation of items to all buyers, i.e., the concatenation of all buyers' allocations.

- (f) The *prices* of items are modeled as $p \in L^1_+$; in other words, the price of item $\theta \in \Theta$ is $p(\theta)$. For the finite case $\Theta = [m]$, we have $p = (p_1, \dots, p_m) \in \mathbb{R}^m_+$.
- (g) For a measurable item subset $A \subseteq \Theta$, let $v_i(A) := \int_A v_i(\theta) d\theta$ (and similarly for p and s), the v_i -induced measure of A . For the finite case $\Theta = [m]$, for any item subset $A \subset [m]$, $v_i(A) = \sum_{j \in A} v_{ij}$ (and similarly for $p(A)$ and $s(A)$).
- (h) Without loss of generality, we assume a unit total budget $\|B\|_1 = 1$, a unit total supply $s(\Theta) = 1$ and normalized buyer valuations $\langle v_i, s \rangle = 1$. In other words, all items have a total value of 1 for every buyer.

Definition 1. Given item prices $p \in L^1_+$, the **demand** of buyer i is its set of utility-maximizing allocations given the prices and budget:

$$D_i(p) := \arg \max \{ \langle v_i, x_i \rangle : x_i \in L^\infty_+, \langle p, x_i \rangle \leq B_i \}.$$

The associated **utility level** $\hat{U}_i(p)$ is defined as the value of $\langle v_i, x_i \rangle$ for any $x_i \in D_i(p)$.

Definition 2. A **market equilibrium (ME)** is an allocation-price pair $(x^*, p^*) \in (L^\infty_+)^n \times L^1_+$ such that the following holds.

- (i) *Supply feasibility*: $\sum_i x_i^* \leq s$.
- (ii) *Buyer optimality*: $x_i^* \in D_i(p^*)$ for all i .

(iii) *Market clearance:* $\langle p^*, s - \sum_i x_i^* \rangle = 0$ (any item with a positive price is fully allocated).

In the above definition and subsequently, all equations involving measurable functions are understood as “holding almost everywhere.” For example, $\sum_i x_i \leq s$ means the (measurable) set $\{\theta \in \Theta : \sum_i x_i(\theta) \leq s(\theta)\}$ has the same measure as Θ . Given a ME (x^*, p^*) , we often denote the (unique) equilibrium utilities as $u_i^* = \langle v_i, x_i^* \rangle$. For a finite-dimensional linear Fisher market, it is well known that a ME can be computed via solving the EG convex program. Recently, [24] generalized this framework to handle the case of an infinite item space. More specifically, consider the following (possibly infinite-dimensional) convex programs.

$$\sup_{x \in (L_+^\infty)^n} \sum_i B_i \log \langle v_i, x_i \rangle \quad \text{s.t.} \quad \sum_i x_i \leq s. \quad (\mathcal{P}_{\text{EG}})$$

$$\inf_{p \in L_+^1, \beta \in \mathbb{R}_+^n} \left(\langle p, s \rangle - \sum_i B_i \log \beta_i \right) \quad \text{s.t.} \quad p \geq \beta_i v_i, \forall i. \quad (\mathcal{D}_{\text{EG}})$$

The following theorem summarizes the results in [24, §3] regarding the above convex programs capturing market equilibria. As shown in that work, the above convex programs satisfy strong duality and their optimal solutions (which correspond to ME) can be characterized by the KKT optimality conditions. We slightly generalize the assumptions of [24] by allowing non-uniform item supplies s instead of $s(\theta) = 1$ for all $\theta \in \Theta$. For completeness, a proof, which is mainly based on the proofs of the results in [24, §3], can be found in the Appendix.

Theorem 1. *The following hold regarding $(\mathcal{P}_{\text{EG}})$ and $(\mathcal{D}_{\text{EG}})$.*

- Both suprema are attained.
- Given x^* feasible to $(\mathcal{P}_{\text{EG}})$ and (p^*, β^*) feasible to $(\mathcal{D}_{\text{EG}})$, they are both optimal if and only if the following holds: (i) $\langle p^*, s - \sum_i x_i^* \rangle = 0$ (market clearance), (ii) $\langle p^* - \beta_i^* v_i, x_i^* \rangle = 0$ (buyer i only receives items within its ‘winning set’ $\{p^* = \beta_i^* v_i\}$) (ii) and $\langle v_i, x_i^* \rangle = u_i^* := B_i / \beta_i^*$ (buyer i gets its maximum possible utility from x_i^*). In this case, (x^*, p^*) is a ME.
- Conversely, for a ME (x^*, p^*) , it holds that (i) x^* is an optimal solution of $(\mathcal{P}_{\text{EG}})$ and (ii) (p^*, β^*) , where $\beta_i^* := B_i / \langle v_i, x_i^* \rangle$, is an optimal solution of $(\mathcal{D}_{\text{EG}})$.

In the above theorem, β_i^* is known as buyer i ’s *utility price*, i.e., price per unit utility at equilibrium. As is well known, in a ME (x^*, p^*) , the allocations x^* are

inline Pareto optimal,

inline envy-free (in a budget weighted sense, i.e., $\langle v_i, x_i^* \rangle / B_i \geq \langle v_i, x_k^* \rangle / B_k$ for all $k \neq i$),

inline proportional (i.e., $\langle v_i, x_i^* \rangle \geq \langle v_i, s \rangle / n = 1/n$); see, e.g., [24, Theorem 3].

Online Fisher markets and equilibria. We now consider a simple online variant of the Fisher market setting, referred to as an online Fisher market (OFM). There are n buyers, each with a valuation $v_i \in L_+^1$. Assume there are discrete time steps $t = 1, 2, \dots$. At each time step t , an item θ_t arrives and each buyer i sees a value $v_i(\theta_t)$. The item must be allocated irrevocably to one buyer. Each buyer i has a budget $B_i > 0$ representing her per-period expenditure rate.¹

Next, we introduce the notions of demand, utility level and online market equilibrium in an OFM. All of them are defined based on sequences of arrived items and their prices; they do not require any distributional assumption on the item arrivals.

Definition 3. *Let the arrived items be $(\theta_\tau)_{\tau \in [t]}$. An allocation (of arrived items) is $(x_i^\tau)_{(\tau, i) \in [t] \times [n]}$, where $x_i^\tau \in [0, 1]$ is the fraction of the item θ_τ allocated to buyer i .² Let the prices of the arrived*

¹This assumption is similar to one made in the literature on budget management in auctions, where each buyers has a per-period expenditure rate and the overall budget equal to the rate times the number of time periods. If a hard budget cap across all time periods is desired, then PACE and similar mechanisms may deplete some buyers’ budgets close to the end of the horizon [6–8].

²We allow fractional allocations in the definition for more generality. As we will see, fractional allocation is not needed: PACE generates allocations and prices that satisfy the OME conditions asymptotically via assigning each arrived item to one buyer.

items be $(p^\tau(\theta_\tau))_{\tau \in [t]}$. The demand of each buyer i (in hindsight) at time t is

$$D_i^t = \arg \max_{(z_i^\tau)_{\tau \in [t]}} \left\{ \frac{1}{t} \sum_{\tau=1}^t v_i(\theta_\tau) z_i^\tau : 0 \leq z_i^\tau \leq 1, \forall \tau, \frac{1}{t} \sum_{\tau=1}^t p^\tau(\theta_\tau) z_i^\tau \leq B_i \right\}. \quad (1)$$

Let \hat{U}_i^t be the utility level associated with this demand, i.e., the maximum value in (1). An **online market equilibrium (OME)** is a pair of allocations $(x_i^\tau)_{(\tau,i) \in [t] \times [n]}$ and prices $p^\tau(\theta_\tau)$ such that the following holds.

- (i) Total allocation does not exceed the unit amount of the item $\sum_i x_i^\tau \leq 1$ for all τ .
- (ii) Buyers' realized allocations are optimal in hindsight: $(x_i^\tau)_{\tau \in [t]} \in D_i^t$ for all i .
- (iii) Market clearance: $\sum_i x_i^\tau = 1$ for τ such that $p^\tau(\theta_\tau) > 0$.

In words, \hat{U}_i^t is the maximum possible (time-averaged) utility buyer i could have attained via choosing from the arrived items $(\theta_\tau)_{\tau \in [t]}$ in hindsight, subject to their respective posted prices $(p^\tau(\theta_\tau))_{\tau \in [t]}$ and her current total budget tB_i , with D_i^t being the set of such utility-maximizing (time-indexed) allocations subject to per-period item availability constraints. An OME is a pair of allocations and prices that make buyers optimal in hindsight and market cleared.

Given an OFM, we define the associated underlying static Fisher market as having the same n buyers and an item space Θ with supply s being the (unknown) distribution from which the arriving items θ_t are drawn. To clarify the concepts of OFM and OME, we consider some simple special cases.

- Suppose all item arrivals $\theta_1, \dots, \theta_t$ are known in advance. Then, the OFM is the same as a static $n \times t$ Fisher market with the same buyers and the t items, each having a unit supply. Here, buyer i 's valuation of item τ is $v_{i\tau} = v_i(\theta_\tau)$. To compute an OME, it suffices to solve the classical (finite-dimensional) Eisenberg-Gale convex program, that is, $(\mathcal{P}_{\text{EG}})$ with $\Theta = [t]$, $s = (1, \dots, 1) \in \mathbb{R}_+^t$ and $x \in \mathbb{R}_+^{n \times t}$. Let the static ME be $(x^*, p^*) \in \mathbb{R}_+^{n \times t} \times \mathbb{R}_{++}^t$. When each item θ_τ arrives, OME allocates a fraction $x_{i\tau}^*$ of the item to each buyer i and set its price as p_τ^* .
- Suppose the sequentially arriving items are drawn i.i.d. from a known underlying distribution $s \in L_+^\infty$ (which specifies a random variable $\theta \sim s$ such that $\mathbf{P}[\theta \in A] = s(A)$ for any measurable set $A \subseteq \Theta$) and all buyers' valuations v_i are known. Suppose we have also computed a static ME (x^*, p^*) (Definition 2) of a market with buyer valuations v_i , budgets B_i and item supplies being the distribution s (the *underlying static market*). Then, when a new item θ_t (which is drawn from the distribution s) arrives at time t , set its price as $p^*(\theta_t)$ and allocate a fraction $x_i^*(\theta_t)/s(\theta_t)$ of it to each buyer i (assume $s(\theta_t) > 0$, i.e., only items with positive supplies can appear). Then, the time-averaged utility of each buyer i is $\frac{1}{t} \sum_{\tau=1}^t v_i(\theta_\tau) x_i^*(\theta_\tau)/s(\theta_\tau)$, which converges to $\mathbf{E}_{\theta \sim s}[v_i(\theta) x_i^*(\theta)/s(\theta)] = \int_\Theta v_i(\theta) x_i^*(\theta) d\theta = u_i^*$ a.s. by the Strong Law of Large Numbers. Since the online process is carried out using static equilibrium prices and allocations, the static ME properties (Definition 2) ensure the required OME properties hold asymptotically.

The above special cases require full knowledge of either the exact future item arrivals or the underlying static market to attain an OME. Next, we propose a simple, distributed dynamics which generates allocations and prices that satisfy the OME conditions asymptotically **without** requiring such knowledge (in particular, without knowledge of the distribution s).

3 The PACE Dynamics

In this section, we introduce the **PACE** (Pacing According to Current Estimated utility) dynamics that prices and allocates sequentially arriving items via (i) maintaining a *pacing multiplier* for each buyer and (ii) simple, distributed updates.³ In §5, we will show that PACE is an instantiation of dual averaging [48], a stochastic first-order method for regularized optimization, applied to a reformulation of $(\mathcal{D}_{\text{EG}})$. In the PACE dynamics, each buyer maintains a pacing multiplier β_i^t , starting from an initial value $\beta_i^1 = 1 + \delta_0$ for some small $\delta_0 > 0$ (e.g., $\delta_0 = 0.05$). At time step t , the following events take place.

³Pacing and pacing multipliers are terminology in budget management in large-scale ad auctions [18, 19].

- (a) An item θ_t appears and each buyer i sees a value $v_i(\theta_t)$ for the item.
- (b) Each buyer i bids their paced value $\beta_i^t v_i(\theta_t)$ for the item.
- (c) The item is allocated to the highest bidder (the *winner* at t): $i_t = \arg \max_i \beta_i^t v_i(\theta_t)$, with ties broken arbitrarily. For concreteness, we always choose the lowest winning index, i.e.,

$$i_t = \min \arg \max_i \beta_i^t v_i(\theta_t).$$

Then, the price of θ_t is set by the first-price rule

$$p^t(\theta_t) = \max_i \beta_i^t v_i(\theta_t) = \beta_{i_t}^t v_{i_t}(\theta_t)$$

and the winner i_t pays this price $p^t(\theta_t)$ for the item θ_t .

- (d) Each buyer i gets a utility

$$u_i^t = v_i(\theta_t) \mathbb{I}\{i = i_t\}.$$

In other words, the winner i_t gets $v_{i_t}(\theta_t)$ and other buyers get zero.

- (e) Each buyer i updates its cumulative average utility \bar{u}_i^t :

$$\bar{u}_i^t = \frac{1}{t} \sum_{\tau=1}^t u_i^\tau = \frac{t-1}{t} \bar{u}_i^{t-1} + \frac{1}{t} u_i^t.$$

- (f) Each buyer i updates their pacing multiplier β_i^{t+1} as follows:

$$\beta_i^{t+1} = \Pi_{[l_i, h_i]}(B_i / \bar{u}_i^t) := \min\{\max\{l_i, B_i / \bar{u}_i^t\}, h_i\}.$$

where $l_i = B_i / (1 + \delta_0)$ and $h_i = 1 + \delta_0$ for some fixed $\delta_0 > 0$ (e.g., $\delta_0 = 0.05$).

As will be seen in §5, buyer i 's equilibrium pacing multiplier (i.e., utility price) satisfies $l_i < \beta_i^* < h_i$ and her per-period utility u_i^t corresponds to the i th component of a stochastic subgradient of a function on β in a reformulation of the convex program (\mathcal{D}_{EG}), on which we run dual averaging. Furthermore, the update rule for β_i^{t+1} is such that, if the realized utilities \bar{u}_i^t were the true static equilibrium utility for buyer i , then β_i^{t+1} would be the equilibrium multiplier. Note that PACE does not randomize (any randomness can only come from the market environment from which item arrivals are drawn) and assigns every item to a single buyer without splitting it.

The simplicity and distributed nature of PACE makes it desirable for large-scale practical use.

- It can be run on arbitrary sequential item arrivals and only requires buyers' valuations $v_i(\theta_t)$ on the arrived items (rather than all valuations v_i over the potentially large item space). No parameter tuning is needed (in particular, no stepsize tuning as in many first-order optimization methods).
- When run as a *centralized* allocation mechanism, PACE only needs to maintain $O(n)$ scalars, namely, β_i^t , B_i and \bar{u}_i^t for all i . At time t , it observes buyers' valuations $v_i(\theta_t)$ of the item θ_t , compute bids $\beta_i^t v_i(\theta_t)$, finds the winner i_t , set the price as the maximal bid $\beta_{i_t}^t v_{i_t}(\theta_t)$ and allocates the item to the winner; finally, it updates \bar{u}^t and β^{t+1} as in (f), which takes $O(n)$ time.
- PACE can also be run among the buyers in a *decentralized* manner, in which case each buyer only maintains two scalar values: the pacing multiplier β_i^t and time-averaged utility \bar{u}_i^t . When a new item arrives, each buyer only performs a few simple arithmetic operations to create a bid $\beta_i^t v_i(\theta_t)$, receives her utility (if she wins) and subsequently updates \bar{u}_i^t and β_i^{t+1} .

These make PACE suitable for Internet-scale online fair division and online Fisher market applications. In particular, it is very reminiscent of how Internet advertising auctions are run. There, a similar auction-based system is used, with the pacing multiplier ensuring that each advertiser smooths out their budget expenditure across the many auctions. The primary difference between this and our setting is that (i) the auction can be first-price or second-price and (ii) buyers usually have *quasilinear* utilities, that is, utility of the item minus the expenditure (price paid) [6–8, 18]. In §C, we extend PACE to quasilinear utilities, which provides a novel online algorithm for first-price pacing equilibrium computation [19].

4 Dual Averaging

In this section, we briefly recap the setup and general convergence results of *dual averaging* [35, 48], which will be used in the analysis of PACE. First, we introduce some notation for this and the next section. Let $\mathbf{e}^{(i)}$ denote the i 'th unit basis vector in \mathbb{R}^n and $\mathbf{1} \in \mathbb{R}^n$ denote the vector of 1's. For $x, y \in \mathbb{R}^n$, $[x, y]$ denotes the Cartesian product of intervals $\prod_{i=1}^n [x_k, y_k] \subseteq \mathbb{R}^n$. All norms $\|\cdot\|$ without a subscript are Euclidean 2-norms, unless otherwise stated.

Let Ψ be a closed, strongly convex function with domain $\text{dom } \Psi := \{w \in \mathbb{R}^n : \Psi(w) < \infty\}$. Here, we do not employ any auxiliary regularizing function, since our problem has a natural source of strong convexity (i.e., a strongly convex Ψ) through the $-B_i \log \beta_i$ terms in $(\mathcal{D}_{\text{EG}})$. Let $Z \subseteq \mathbb{R}^d$ be an arbitrary sample space. For each $z \in Z$, let f_z be a convex and subdifferentiable function on $\text{dom } \Psi$. Consider the following regularized convex optimization problem [48, §1.1]:

$$\min_w \mathbf{E} f_z(w) + \Psi(w), \quad (2)$$

where the expectation is taken over a probability distribution \mathcal{D} on Z . We assume access to an oracle that, given any f_t and $w \in \text{dom } \Psi$, returns a subgradient $g^t \in \partial f_t(w)$. The dual averaging algorithm (DA) [48, Algorithm 1], with a strongly convex Ψ and no auxiliary regularizer, is as follows. First, set $w_1 \in \arg \min_w \Psi(w)$ and $\bar{g}^0 = 0$. Then, for each $t = 1, 2, \dots$, DA performs the following steps:

- (1) Observe f_t and compute $g^t \in \partial f_t(w^t)$.
- (2) Update the average subgradient (the *dual average*) via $\bar{g}^t = \frac{t-1}{t} \bar{g}^{t-1} + \frac{1}{t} g^t$.
- (3) Compute the next iterate $w^{t+1} = \arg \min_w \{\langle \bar{g}^t, w \rangle + \Psi(w)\}$.

The following convergence guarantee on DA is proved as part of the proof of Corollary 4 in [48].

Theorem 2. *Dual averaging generates iterates w^t such that*

$$\mathbf{E} \|w^t - w^*\|^2 \leq \frac{(6 + \log t) G^2}{t \sigma^2},$$

where G^2 is an upper bound on $\mathbf{E} \|g^t\|^2$, $t = 1, 2, \dots$ and σ is the strong convexity modulus of Ψ .

When solving the stochastic optimization problem (2), in Theorem 2, we can set G^2 to be an upper bound on $\sup_{w \in \text{dom } \Psi} \mathbf{E} \|g_z(w)\|^2$, where $g_z(w)$ is a subgradient oracle mapping each $(z, w) \in Z \times \text{dom } \Psi$ to a subgradient and the expectation is over $z \sim \mathcal{D}$ and possible randomness of the subgradient oracle. We will shortly see that a reformulation of \mathcal{D}_{EG} , when cast into the form (2), exhibits stochastic subgradients that are exactly buyers' received utilities in each time step. Using Theorem 2, we can show that the sequence of pacing multipliers β^t generated by PACE converges to the underlying (equilibrium) utility prices β^* of the static Fisher market.

5 Convergence Analysis of the PACE dynamics

We will now show that PACE correspond to running DA on the vector β^t of pacing multipliers for the buyers. To this end, we first reformulate $(\mathcal{D}_{\text{EG}})$ into a (finite-dimensional) convex program in β in the form of (2):

$$\min_{\beta} \left(\left\langle \max_i \beta_i v_i, s \right\rangle - \sum_i B_i \log \beta_i \right) \text{ s.t. } \beta \in [B/(1 + \delta_0), (1 + \delta_0)\mathbf{1}], \quad (3)$$

where $\delta_0 > 0$ is an arbitrarily small constant. The bounds on β do not change the optimal solution, because $\beta_i^* \in (B_i, 1)$ for each i . Detailed steps of the reformulation are given in Appendix B.

In order to run DA, we need to compute a subgradient of $f_{\theta} : \beta \mapsto \max_i \beta_i v_i(\theta)$ at any $\theta \in \Theta$. Following [24, §5], since f_{θ} is a piecewise linear function, a subgradient is

$$g_{\theta}(\beta) := v_{i^*}(\theta) \mathbf{e}^{(i^*)} \in \partial f_{\theta}(\beta),$$

where $i^* = \min \arg \max_i \beta_i v_i(\theta)$ is the winner (see, e.g., [10, Theorem 3.50]).

We can now show that the PACE dynamics corresponds to running DA on (3). Here, $\Psi(\beta) = -\sum_i B_i \log \beta_i$ with $\text{dom } \Psi = [B/(1 + \delta_0), (1 + \delta_0)\mathbf{1}]$. First, choose $\beta^1 = \arg \min \Psi = (1 + \delta_0)\mathbf{1}$ (i.e., $\beta_i^1 = 1 + \delta_0$ for all i) and $\bar{g}^0 = 0$. At each time step $t = 1, 2, \dots$, given the current pacing multiplier β^t , DA applied to (3) unrolls the following steps.

- An item θ_t arrives, having value $v_i(\theta_t)$ for each buyer i . The function f_t in DA is

$$f_{\theta_t} : \beta \mapsto \max_i \beta_i v_i(\theta_t).$$

- The winner is $i_t = \min \arg \max_i \beta_i^t v_i(\theta_t)$ and a subgradient is $g^t = v_{i_t}(\theta_t) \mathbf{e}^{(i_t)} \in \partial f_t(\beta^t)$. Its i th entry is exactly the realized (single-period) utility of individual i at time t in PACE, that is, $g_i^t = v_i(\theta_t) \mathbb{I}\{i = i_t\} = u_i^t$.
- Update the dual average (time-averaged utilities): for each i , compute $\bar{g}^t = \frac{t-1}{t} \bar{g}^{t-1} + \frac{1}{t} g^t$, i.e.,

$$\bar{g}_i^t = \frac{t-1}{t} \bar{g}_i^{t-1} + \frac{1}{t} v_i(\theta_t) \mathbb{I}\{i = i_t\}.$$

- Update the pacing multipliers:

$$\beta^{t+1} = \arg \min_{\beta \in [B/(1+\delta_0), (1+\delta_0)\mathbf{1}]} \left\{ \langle \bar{g}^t, \beta \rangle - \sum_i B_i \log \beta_i \right\}.$$

The minimization problem is separable in each i and exhibits a simple and explicit solution which recovers step (f) in PACE (where $\bar{g}_i^t = \bar{u}_i^t$):

$$\beta_i^{t+1} = \arg \min_{\beta_i \in [B/(1+\delta_0), 1+\delta_0]} \{ \bar{g}_i^t \beta_i - B_i \log \beta_i \} \Rightarrow \beta_i^{t+1} = \Pi_{[B/(1+\delta_0), 1+\delta_0]} \left(\frac{B_i}{\bar{u}_i^t} \right).$$

As mentioned earlier, PACE does not require a stepsize parameter. This is because DA is stepsize-free given a strongly convex regularizer Ψ , which is indeed the case in our reformulation (3). In addition, in the above update step for β_i^{t+1} , the directions of change are as follows.

- For a non-winner $i \neq i_t$, we have $u_i^t = 0$ and hence $\bar{u}_i^t \leq \bar{u}_i^{t-1}$. This implies $\beta_i^{t+1} \geq \beta_i^t$. In words, a non-winner's pacing multiplier weakly increases. The increase is strict if $\bar{u}_i^{t-1} > 0$, i.e., buyer i has already received a nonzero utility.
- For the winner i_t , $\bar{u}_{i_t}^t$ may become greater than $\bar{g}_{i_t}^{t-1}$, in which case $\beta_{i_t}^{t+1} \leq \beta_{i_t}^t$. In words, the winner's pacing multiplier may go up or down.

In order to analyze PACE, we assume $v_i(\Theta) = 1$, $v_i \in L_+^\infty$ (normalized and a.e.-bounded valuations)⁴ and that there is an underlying item distribution $s \in L_+^\infty$ from which the item arrivals θ_t , $t = 1, 2, \dots$ are drawn i.i.d.⁵ Define the underlying static Fisher market as one having the same n buyers (each with valuation v_i and budget B_i) and item supplies s . Denote the equilibrium utilities and utility prices w.r.t. the underlying static market as u^* and β^* , respectively. We further assume that the valuations are $v_i \in L_+^\infty$ (i.e., a.e.-bounded on the item space). This is not restrictive: since an individual item θ has value $v_i(\theta)$ for each buyer i , it should be a finite value.

Convergence of pacing multipliers. After aligning PACE with DA, the convergence of the pacing multipliers β^t follows directly from Theorem 2.

Theorem 3. *PACE generates pacing multipliers β^t such that*

$$\mathbf{E} \|\beta^t - \beta^*\|^2 \leq \frac{(6 + \log t) G^2}{t \sigma^2}, \quad t = 1, 2, \dots,$$

where $G^2 = \max_i \mathbf{E}_{\theta \sim s} [v_i(\theta)^2] \leq \max_i \|v_i\|_\infty^2$, $\sigma = \frac{\min_i B_i}{(1+\delta_0)^2}$.

In other words, we have mean-square convergence of β^t to β^* at a $O((\log t)/t)$ rate. Since $\|B\|_1 = 1$, we have $\min_i B_i \leq 1/n$. Hence, $\sigma = O(1/n)$ and the constant in the bound is $\Omega(n^2)$. Whether such dependence on n can be improved via new analysis remains an interesting research question.

Convergence of utilities. We next show that the time-averaged utility \bar{u}^t (which equals to the dual average \bar{g}^t) converges to the equilibrium utility vector u^* of the underlying Fisher market. A key step in the proof is to bound the probability of a projection in updating β_i^{t+1} , that is, $\mathbf{P}[B_i/\bar{u}_i^t \notin [l_i, u_i]]$.

⁴The a.e.-boundedness assumption is needed in subsequent convergence analysis. Since Θ has a finite measure, it holds that $L_+^\infty \subseteq L_+^1$. For a finite item space $\Theta = [m]$, both are equal to \mathbb{R}_+^m .

⁵The distributional assumption on item arrivals (i.e., they are drawn i.i.d. from an unknown distribution s) is needed to establish asymptotic equilibrium properties of PACE. See Appendix B for an example that any algorithm can yield arbitrarily suboptimal allocations without such a distributional assumption.

Theorem 4. For each i , let $\epsilon_i := \min\{h_i - \beta_i^*, \beta_i^* - l_i\} > 0$ be the minimum distance to the pacing-multiplier interval. Let $\|v\|_\infty := \max_i \|v_i\|_\infty$. It holds that

$$\mathbf{E}(\bar{u}_i^t - u_i^*)^2 \leq \left(\frac{\|v_i\|_\infty^2}{\epsilon_i^2} + \left(\frac{1 + \delta_0}{B_i} \right)^2 \right) \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2.$$

Hence, letting $C = \frac{1}{(\min_i B_i)^2} ((\|v\|_\infty/\delta_0)^2 + (1 + \delta_0)^2)$, we have

$$\mathbf{E}\|\bar{u}^t - u^*\|^2 \leq C \cdot \frac{(6 + \log(t+1))G^2}{(t+1)\sigma^2}.$$

Note that $C = \Omega(n^2)$. Hence, in this and the next theorems, the constant in the bound is $\Omega(n^4)$, which arises from C and $\sigma = O(\min_i B_i) = O(1/n)$.

Convergence of expenditures. The expenditure of buyer i at time step t is $b_i^t = \beta_i^t v_i(\theta_t) \mathbb{I}\{i = i_t\}$. In other words, only the winner i_t spends a nonzero amount, which is its bid. Let $\bar{b}_i^t = \frac{1}{t} \sum_{\tau=1}^t b_i^\tau$ be buyer i 's average expenditure. Utilizing the above convergence results, we show mean-squared convergence of \bar{b}^t to B at an $O((\log t)^2/t)$ rate.

Theorem 5. For each i , it holds that

$$\mathbf{E}(\bar{b}_i^t - B_i) \leq 2 \left[(\beta_i^*)^2 \mathbf{E}(\bar{g}_i^t - u_i^*)^2 + 2\|v_i\|_\infty^2 \frac{1}{t} \sum_{\tau=1}^t \mathbf{E}(\beta_i^\tau - \beta_i^*)^2 \right].$$

For $t \geq 3$ and the constant C defined in Theorem 4, we have

$$\mathbf{E}\|\bar{b}^t - B\|^2 \leq \frac{2G^2}{t\sigma^2} \left(6(C + \|v\|_\infty^2) + (C + 6\|v\|_\infty^2) \log t + \frac{\|v\|_\infty^2}{2} (\log t)^2 \right).$$

PACE attains OME asymptotically. Next, we show that PACE attains OME asymptotically, i.e., it generates allocations and prices that make buyers *no-regret* and *envy-free* in the limit (these notions will be clarified shortly). Let $x_i^t := \mathbb{I}\{i = i_t\}$ denote whether buyer i_t is the winner (i.e., whether she is allocated the item θ_t at time t). Utilizing Theorems 4 and 5, we can show that buyer i 's regret, that is, the difference between the maximum possible utility in hindsight \hat{U}_i^t (Definition 3) and the realized utility \bar{u}_i^t , vanishes as t grows. The same holds for each buyer's envy. In other words, at a large t , in hindsight, no buyer prefers another buyer's set of allocated items (up to a vanishing error).⁶

Theorem 6. Denote $\xi_i^t = |\bar{u}_i^t - u_i^*|$, $\Delta_i^t = |\bar{b}_i^t - B_i|$, $\gamma_t = \frac{\|v\|_\infty}{t} \sum_{\tau=1}^t \|\beta^\tau - \beta^*\|_\infty$. Let $r_i^t := \max\{\hat{U}_i^t - \bar{u}_i^t, 0\}$ be the regret of buyer i at time t . Then, it holds that $r_i^t \leq \xi_i^t + \gamma_t/B_i$ and $\mathbf{E}(r_i^t)^2 = O((\log t)^2/t)$. Furthermore, define the envy of buyer i (w.r.t. all other buyers) at time t be $\rho_i^t = \max_k \bar{u}_{ik}^t/B_k - \bar{u}_i^t/B_i$, where $\bar{u}_{ik}^t = \frac{1}{t} \sum_{\tau=1}^t v_i(\theta_\tau) x_k^\tau$ is buyer i 's time-averaged utility given her own valuations and of buyer k 's allocations. Denote $\eta_i^t = \frac{1}{t} \sum_{\tau=1}^t (p^*(\theta_t) - \beta_i^\tau v_i(\theta_t)) x_i^\tau$. It holds that

$$\rho_i^t \leq \frac{1}{B_i} \left(\xi_i^t + \max_{k \neq i} \frac{\Delta_k^t + \eta_k^t}{B_k} \right) \text{ and } \mathbf{E}(\eta_i^t)^2 \leq \frac{\|v\|_\infty^2 G^2}{t\sigma^2} \left(6(1 + \log t) + \frac{(\log t)^2}{2} \right).$$

Hence, the envy ρ_i^t of buyer i vanishes in mean square, i.e., $\mathbf{E}(\rho_i^t)^2 = O((\log t)^2/t)$.

In light of Definition 3, Theorem 6 shows that $(x_i^\tau)_{(i,\tau) \in [n] \times [t]}$ is approximately optimal for buyer i . Since PACE also clears the market, we conclude that it attains OME asymptotically. Recall that theorem 4 ensures that buyers' \bar{u}_i^t converge to their static equilibrium utilities u_i^* . Since the latter satisfy Pareto optimality and proportional share guarantee ($u_i^* \geq B_i$ for all i), so are the time-averaged realized utilities in the limit. Together with Theorem 6, we conclude that PACE achieves the said fairness and efficiency guarantees, namely, Pareto optimality, envy-freeness and proportional-share guarantee, asymptotically.

⁶In a static market, given an allocation $x \in \mathbb{R}_+^{n \times m}$, the (maximum, budget-weighted) envy of buyer i toward others' bundles is $\rho_i(x) = \max_k \langle v_i, x_k \rangle / B_k - \langle v_i, x_i \rangle / B_i$ (see, e.g., [12, 46]). It is well-known that $\rho_i(x^*) = 0$ for all i at equilibrium, a consequence of buyer optimality (Definition 2).

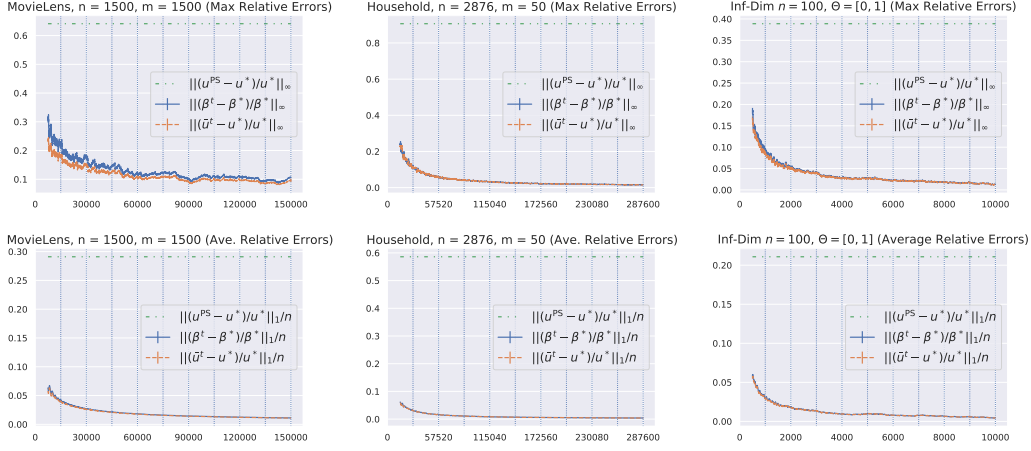


Figure 1: In all of our markets, iterates of the PACE dynamics quickly converges to their static equilibrium values both in the average case and the worst-off-buyer case. The horizontal line shows the fraction of u^* achieved by the proportional share solution. The PACE utilities quickly outperform the proportional share utilities. Vertical lines indicate when t is a multiple of $10n$.

6 Experiments

We evaluate the PACE dynamics in several real and synthetic datasets, namely, MovieLens, Household Items and an infinite-dimensional market instance with item space $\Theta = [0, 1]$ and v_i being linear functions on $[0, 1]$. For the first two datasets, see [31] for more information and exploratory data analysis. For all datasets, we consider the CEEI (fair division) setting where $B_i = 1/n$ for all i . For each dataset (with number of buyers $n = 1500, 2876, 100$, respectively), we run PACE for $T = 10n$ time steps (iterations). More details on the experiments and additional plots displaying convergence of expenditures can be found in Appendix D. Figure 1 displays the mean values of the average and maximum relative errors of the pacing multipliers and time-averaged cumulative utilities over 10 repeated experiments with different seeds (relative errors of cumulative spending w.r.t. total budgets are plotted separately in Appendix D). The standard errors are also displayed as vertical bars but are very small and nearly invisible. Vertical dotted lines indicate $t = 10n$. The figures do not show the initial iterates $t = 1, \dots, 5n$.

We see that PACE converges very quickly numerically: within 10 epochs ($10n$ time steps) average deviations in most quantities falls within 5% of the equilibrium quantity, with the worst case not far behind. An important point is that budget spend takes much longer to converge than utility. This demonstrates an important practical difference for using PACE in an allocation scenario where budgets are ‘real money’ (e.g. Internet ad impressions) as compared to a CEEI-like setting, where budgets are faux currency only used for fair division.

7 Conclusion

We introduced the concept of an online Fisher market and proposed the PACE dynamics. We showed that when items arrive sequentially and stochastically, PACE converges to equilibrium outcomes of the underlying market model. Furthermore, we showed that, as a consequence of this, PACE can be used in online fair division problems to generate an online allocation that, asymptotically, achieves the compelling fairness properties of CEEI.

Many questions remain for future research. We mostly focused on the case where budgets are faux currency and there are many open questions for adapting PACE to a real-money budget-management setting as well as more complicated nonlinear utility models. Another imperative question, especially for practitioners, is whether PACE guarantees some level of incentive-compatibility.

References

- [1] M. Aleksandrov and T. Walsh. Online fair division: A survey. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13557–13562, 2020.
- [2] M. Aleksandrov, H. Aziz, S. Gaspers, and T. Walsh. Online fair division: Analysing a food bank problem. *arXiv preprint arXiv:1502.07571*, 2015.
- [3] C. Arnsperger. Envy-freeness and distributive justice. *Journal of Economic Surveys*, 8(2): 155–186, 1994.
- [4] Y. Azar, N. Buchbinder, and K. Jain. How to allocate goods in an online market? *Algorithmica*, 74(2):589–601, 2016.
- [5] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015.
- [6] S. Balseiro, A. Kim, M. Mahdian, and V. Mirrokni. Budget management strategies in repeated auctions. In *Proceedings of the 26th International Conference on World Wide Web*, pages 15–23, 2017.
- [7] S. R. Balseiro and Y. Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 65(9):3952–3968, 2019.
- [8] S. R. Balseiro, O. Besbes, and G. Y. Weintraub. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science*, 61(4):864–884, 2015.
- [9] M. H. Bateni, Y. Chen, D. Ciocan, and V. Mirrokni. Fair resource allocation in a volatile marketplace. *Available at SSRN 2789380*, 2018.
- [10] A. Beck. *First-order methods in optimization*, volume 25. SIAM, 2017.
- [11] B. Birnbaum, N. R. Devanur, and L. Xiao. Distributed algorithms via gradient descent for fisher markets. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 127–136. ACM, 2011.
- [12] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [13] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 305–322. ACM, 2016.
- [14] X. Chen and S.-H. Teng. Spending is not easier than trading: on the computational equivalence of fisher and arrow-debreu equilibria. In *International Symposium on Algorithms and Computation*, pages 647–656. Springer, 2009.
- [15] Y. K. Cheung, M. Hoefer, and P. Nakhe. Tracing equilibrium in dynamic markets via distributed adaptation. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1225–1233, 2019.
- [16] R. Cole and V. Gkatzelis. Approximating the Nash social welfare with indivisible items. *SIAM Journal on Computing*, 47(3):1211–1236, 2018.
- [17] R. Cole, N. R. Devanur, V. Gkatzelis, K. Jain, T. Mai, V. V. Vazirani, and S. Yazdanbod. Convex program duality, fisher markets, and Nash social welfare. In *18th ACM Conference on Economics and Computation, EC 2017*. Association for Computing Machinery, Inc, 2017.
- [18] V. Conitzer, C. Kroer, E. Sodomka, and N. E. Stier-Moses. Multiplicative pacing equilibria in auction markets. In *International Conference on Web and Internet Economics*, 2018.
- [19] V. Conitzer, C. Kroer, D. Panigrahi, O. Schrijvers, E. Sodomka, N. E. Stier-Moses, and C. Wilkens. Pacing equilibrium in first-price auction markets. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. ACM, 2019.

- [20] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- [21] A. Dvoretzky, A. Wald, J. Wolfowitz, et al. Relations among certain ranges of vector measures. *Pacific Journal of Mathematics*, 1(1):59–74, 1951.
- [22] E. Eisenberg and D. Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.
- [23] Y. Gao and C. Kroer. First-order methods for large-scale market equilibrium computation. In *Neural Information Processing Systems 2020, NeurIPS 2020*, 2020.
- [24] Y. Gao and C. Kroer. Infinite-dimensional fisher markets and tractable fair division. *arXiv preprint arXiv:2010.03025*, 2021.
- [25] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *Nsdi*, volume 11, pages 24–24, 2011.
- [26] K. Goldberg, T. Roeder, D. Gupta, and C. Perkins. Eigentaste: A constant time collaborative filtering algorithm. *information retrieval*, 4(2):133–151, 2001.
- [27] F. M. Harper and J. A. Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.
- [28] L. Kantorovich. Mathematics in economics: achievements, difficulties, perspectives. Technical report, Nobel Prize Committee, 1975.
- [29] I. Kash, A. D. Procaccia, and N. Shah. No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research*, 51:579–603, 2014.
- [30] C. Kroer and A. Peysakhovich. Scalable fair division for ‘at most one’ preferences. *arXiv preprint arXiv:1909.10925*, 2019.
- [31] C. Kroer, A. Peysakhovich, E. Sodomka, and N. E. Stier-Moses. Computing large market equilibria using abstractions. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 745–746, 2019.
- [32] D. C. McElfresh, C. Kroer, S. Pupyrev, E. Sodomka, K. A. Sankararaman, Z. Chauvin, N. Dexter, and J. P. Dickerson. Matching algorithms for blood donation. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 463–464, 2020.
- [33] R. Murray, C. Kroer, A. Peysakhovich, and P. Shah. Robust market equilibria with uncertain preferences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2192–2199, 2020.
- [34] R. Murray, C. Kroer, A. Peysakhovich, and P. Shah. <https://research.fb.com/blog/2020/09/robust-market-equilibria-how-to-model-uncertain-buyer-preferences/>, Sep 2020.
- [35] Y. Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical programming*, 120(1):221–259, 2009.
- [36] Y. Nesterov and V. Shikhman. Computation of fisher–gale equilibrium by auction. *Journal of the Operations Research Society of China*, 6(3):349–389, 2018.
- [37] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic game theory*. Cambridge University Press, 2007.
- [38] A. Othman, T. Sandholm, and E. Budish. Finding approximate competitive equilibria: efficient and fair course allocation. In *AAMAS*, volume 10, pages 873–880, 2010.
- [39] A. Othman, C. Papadimitriou, and A. Rubinstein. The complexity of fairness through equilibrium. *ACM Transactions on Economics and Computation (TEAC)*, 4(4):1–19, 2016.

- [40] D. C. Parkes, A. D. Procaccia, and N. Shah. Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. *ACM Transactions on Economics and Computation (TEAC)*, 3(1):1–22, 2015.
- [41] A. Peysakhovich and C. Kroer. Fair division without disparate impact. *Mechanism Design for Social Good*, 2019.
- [42] B. Plaut and T. Roughgarden. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics*, 34(2):1039–1068, 2020.
- [43] H. Scarf et al. *On the computation of equilibrium prices*. Number 232. Cowles Foundation for Research in Economics at Yale University New Haven, CT, 1967.
- [44] V. I. Shmyrev. An algorithm for finding equilibrium in the linear exchange model with fixed budgets. *Journal of Applied and Industrial Mathematics*, 3(4):505, 2009.
- [45] S. R. Sinclair, S. Banerjee, and C. L. Yu. Sequential fair allocation: Achieving the optimal envy-efficiency tradeoff curve. *arXiv preprint arXiv:2105.05308*, 2021.
- [46] H. R. Varian et al. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974.
- [47] V. V. Vazirani and M. Yannakakis. Market equilibrium under separable, piecewise-linear, concave utilities. *Journal of the ACM (JACM)*, 58(3):1–25, 2011.
- [48] L. Xiao. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 11:2543–2596, 2010.
- [49] L. Zhang. Proportional response dynamics in the fisher market. *Theoretical Computer Science*, 412(24):2691–2698, 2011.

Checklist. We answered the questions as accurately as possible and believe that many of them do not apply to our work. We are happy to elaborate further should there be any questions or concerns.

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? Yes.
 - (b) Did you describe the limitations of your work? Yes.
 - (c) Did you discuss any potential negative societal impacts of your work? N.A.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? Yes.
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? Yes.
 - (b) Did you include complete proofs of all theoretical results? Yes.
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.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? N.A.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? Yes.
 - (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)? No. (N.A.)
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? Yes.
 - (b) Did you mention the license of the assets? N.A.
 - (c) Did you include any new assets either in the supplemental material or as a URL? No.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? Yes.
 - (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? N.A.
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? N.A.
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? N.A.

A Related Work

The problem of market equilibrium computation has been of interest in economics for a long time (see, e.g., [37]). There is a large literature focusing on computation of equilibrium in the specific case of (finite-dimensional) Fisher markets through various convex optimization formulations [17, 22, 31, 44] and gradient-based methods [11, 23, 36]. Other works extend these results to settings such as quasilinear utilities, capped utilities, indivisible items, or imperfectly specified utility functions [13, 16, 17, 30, 33, 41]. One of the most well-known algorithm for computing static market equilibria under *Constant Elasticity of Substitution* (CES) utilities is the *Proportional Response* (PR) dynamics [11, 49]. Recently, [24] extends the classical Fisher market model to a measurable (possibly continuous) item space and shows that infinite-dimensional EG-type convex programs capture ME under this setting. Our work extends these ideas to a Fisher market-like scenario where items arrive sequentially.

The Fisher market literature above focuses on divisible items or randomized allocations of indivisible items. There is also a large literature on fair allocation of indivisible items (e.g. [5, 13, 42]) including approximate ME-based methods [12, 39]. We note that all allocations in our setting are discrete and the relationship to Fisher markets happens in the time-average sense.

Perhaps most similar to our setting is that of [4], who study how to allocate items in an online fashion in order to obtain a market-equilibrium-like allocation. However, they consider competitive ratios, and give a primal-dual algorithm that suffers at most a logarithmic loss compared to the best hindsight optimal solution, even for worst-case arrivals. In addition to the lack of asymptotic convergence, they also only show guarantees on various (arithmetic, geometric, harmonic) averages of the utilities. In contrast to this, our work considers stochastic arrivals, and gives an adaptive algorithm for asymptotically achieving all the desirable market equilibrium properties (e.g. no envy, Pareto optimality, equilibrium utilities). Another important difference is that our approach is easily implemented as a distributed dynamics that requires only a first-price auction allocation mechanism with indivisible allocations, in $O(n)$ time per item arrival (n being the number of buyers). At each time step, the algorithm only uses buyers' valuations of the current arriving item. This makes our approach suitable for implementation in large-scale systems with a huge, possibly infinite item space.

Other methods for online fair division have been studied by various authors. In this literature, there are various notions of "online:" either buyers, items, or both can arrive online. Here we survey only related work where items arrive online. [2] studies a simple mechanism where agents can declare if they like an item, and then a coin is flipped to determine which of the agents that liked the item will get it. [29] studies online allocation for *Leontief* utilities (where each agent wants a bundle with items of fixed proportions) and shows how to achieve various properties for this setting. [15] studies an evolving market environment and shows that the PR dynamics generates iterates that are close to the (changing) equilibrium. Similar to the classical PR dynamics, in each time step, all buyers' valuations of all items are known to the algorithm. In contrast, our work allows an *unknown* underlying market from which items are sampled: in each time step, PACE only uses buyers' valuations of the current arriving item. [9] studies an online fair allocation problem and proposes a stochastic approximation scheme, which relies on frequently resolving the EG convex program, that ensures a constant approximation ratio (in terms of a proportional fairness measure) relative to the offline fair allocation. Very recently, [45] studies the problem of online fair division in which there is a fixed (finite) number of divisible items and sequentially arriving agents. The authors show that envy-freeness and (Pareto) efficiency cannot be minimized simultaneously; instead, there exists a boundary such that any algorithm can only possibly achieve the envy-efficiency combinations on one side of it, while they also propose such an algorithm. See also [1] for a survey of further works in this area.

The idea of pacing has been studied in the context of budget management in second-price auctions. The work most related to our work is [7], which studies an online version of that setting. [7] and our work are very different in terms of problem settings, results, and analysis. Here, we point out some key differences. [7] first show that, in the typical setting of a single bidder interacting with second-price auctions where values and prices are drawn from a stochastic environment, an adaptive pacing strategy achieves $O(1/\sqrt{T})$ regret and ergodic convergence of the bidder's pacing multipliers. Then, under additional assumptions on monotonicity of the bidders' expected expenditures, the authors establish *game-theoretic equilibrium properties* when all bidders use the same strategy, i.e., under the "simultaneous learning" setting. In contrast, we focus on *market equilibrium properties* such as fairness and efficiency. Furthermore, in terms of technical assumptions and convergence

results, our PACE dynamics is stepsize-free, both in theory and numerically, while the adaptive pacing algorithm in [7] requires careful stepsizing rules; in our setting, we have last-iterate convergence of pacing multipliers (Theorem 7), whereas [7] only establishes ergodic convergence without a rate. These differences are, fundamentally, due to the use of different first-order optimization methods, and the fact that we leverage strong convexity of the EG dual. The analysis in [7] builds upon the convergence properties of mirror descent and stochastic approximation. Hence, it requires pre-determined vanishing stepsizes (their Assumption 1). In the simultaneous learning setting, the stepsizes of all bidders furthermore need to be carefully selected in joint fashion (their Assumption 3). In contrast, PACE is completely stepsize free, thanks to structure of the reformulated dual EG convex program (3), on which dual averaging can be applied directly without any stepsize parameter or an auxiliary regularization term. This makes it vastly easier to apply PACE in practice.

B Proofs, Derivations and Examples

Proof of Theorem 1

In [24, §3], it is assumed that item supplies are uniform, i.e., $s(\theta) = 1$ for all $\theta \in \Theta$. As is well-known in the finite case ($\Theta = [m]$), this assumption is w.l.o.g. when studying a static Fisher market. Here, we show that all results in [24, §3] can be easily generalized to the case of non-uniform supplies s (Theorem 1). For any market instance M with buyer valuations $v_i \in L_+^1$, budgets B_i , $i \in [n]$ and item supplies $s \in L_+^\infty$ (all normalized as described in (h) in §2), consider another market instance \tilde{M} with supplies being the constant function taking 1 on Θ (denoted as $\mathbf{1}$), valuations

$$\tilde{v}_i(\theta) = v_i(\theta)s(\theta)$$

and the same budgets. First, note that $\tilde{v}_i \in L_+^1$ since $s \in L_+^\infty$ and Θ has a finite measure ($\|s\|_\infty = \inf\{M : |s| \leq M \text{ a.e.}\}$):

$$\int_{\Theta} \tilde{v}_i(\theta) d\theta \leq \|s\|_\infty \int_{\Theta} v_i(\theta) d\theta = \|s\|_\infty v_i(\Theta) = \|s\|_\infty.$$

Denote the set of feasible allocations of M as F , that is, the set of (x_i) such that $x_i \in L_+^\infty$ for all i and $\sum_i x_i \leq s$. Similarly, denote the set of feasible allocations of \tilde{M} as \tilde{F} . For any $x \in F$, consider

$$\tilde{x}_i(\theta) = \begin{cases} x_i(\theta)/s(\theta) & \text{if } \theta > 0, \\ 0 & \text{o.w.} \end{cases} \quad (4)$$

Since $0 \leq x_i \leq s$ (a.e.), we have $0 \leq \tilde{x}_i \leq \mathbf{1}$ (which means $\tilde{x}_i \in L_+^\infty$). Since $\sum_i x_i \leq s$, we have

$$\sum_i \tilde{x}_i \leq \mathbf{1}.$$

Therefore, the set of utilities attainable by allocations in F is the same as the set of utilities attainable by \tilde{F} . In other words,

$$U = \{u \in \mathbb{R}_+^n : x \in F, \langle v_i, x_i \rangle = u_i, i \in [n]\} = \tilde{U} = \{\tilde{u} \in \mathbb{R}_+^n : \tilde{x} \in \tilde{F}, \langle \tilde{v}_i, \tilde{x}_i \rangle = \tilde{u}_i\}.$$

By [24, Lemma 1] and its proof there (here, we only need the compactness, not the existence of a pure allocation for any feasible utility vector; hence, invoking [21, Theorem 1] suffices), \tilde{U} is convex and compact and so is U . Hence, the suprema of $(\mathcal{P}_{\text{EG}})$ is attained. Completely analogous to the proof of [24, Theorem 2], we can show that both suprema of $(\mathcal{P}_{\text{EG}})$ is attained. Furthermore, completely analogous to the proofs of Lemma 3 and Theorem 2 there, we can show that strong duality holds for $(\mathcal{P}_{\text{EG}})$ and $(\mathcal{D}_{\text{EG}})$. More specifically, for any x feasible to $(\mathcal{P}_{\text{EG}})$ and (p, β) feasible to $(\mathcal{D}_{\text{EG}})$, and

any $0 \leq u_i \leq \langle v_i, x_i \rangle$ (i.e., adding auxiliary variable u_i to $(\mathcal{P}_{\text{EG}})$), it holds that

$$\begin{aligned}
\sum_i B_i \log u_i &\leq \sum_i B_i \log u_i - \sum_i \beta_i (u_i - \langle v_i, x_i \rangle) - \left\langle p, \sum_i x_i - s \right\rangle \\
&= \sum_i (B_i \log u_i - \beta_i u_i) - \sum_i \langle p - \beta_i v_i, x_i \rangle + \langle p, s \rangle \\
&\leq \sum_i (B_i \log \frac{B_i}{\beta_i} - \beta_i \cdot \frac{B_i}{\beta_i}) - \sum_i \langle p - \beta_i v_i, x_i \rangle + \langle p, s \rangle \\
&\leq \sum_i (B_i \log B_i - B_i) + \langle p, s \rangle - \sum_i B_i \log \beta_i \\
&= \langle p, s \rangle - \sum_i B_i \log \beta_i - C,
\end{aligned}$$

where the constant $C = \sum_i B_i (1 - \log B_i)$. In the above derivation, the first inequality uses $\beta_i \geq 0$, $u_i \leq \langle v_i, x_i \rangle$, $\sum_i x_i \leq s$; the second inequality uses the fact that $u_i = B_i/\beta_i$ maximizes the function

$$u_i \mapsto B_i \log u_i - \beta_i u_i$$

for any $\beta_i > 0$ (i.e., substituting $u_i = B_i/\beta_i$ into the first line); the third inequality uses feasibility w.r.t. $(\mathcal{D}_{\text{EG}})$, i.e., $p \geq \beta_i v_i$ for all i . Hence, when all inequalities are tight at a pair of solutions x^* and (p^*, β^*) feasible to $(\mathcal{P}_{\text{EG}})$ and $(\mathcal{D}_{\text{EG}})$, respectively (i.e., both optima are attained), the following KKT conditions must hold:

- $\langle p^*, s - \sum_i x_i^* \rangle = 0$ (via the first inequality above being tight).
- $u^* = B_i/\beta_i^*$ for all i (via the second above).
- $\langle p^* - \beta_i^* v_i, x_i^* \rangle = 0$ for all i (via the third above).

As the proof of [24, Theorem 2] shows, these conditions (together with feasibility w.r.t. the two convex programs) are necessary and sufficient for (x^*, p^*) being a ME.

An example: necessity of the distributional assumption on item arrivals

In the definitions of OFM and OME in §2, we do not impose any distributional assumption on the sequentially arriving items θ_t . The PACE dynamics does not require any distributional assumption either. It is in the analysis of PACE in §5 that we assume that θ_t are drawn i.i.d. from an (unknown) underlying distribution $s \in L_+^\infty$ (where $s(\Theta) = 1$ since it is a distribution). We define an underlying static Fisher market with the same buyers and item supplies s . Then, §5 essentially shows that PACE guarantees that various (time-averaged) quantities converge to their static equilibrium quantities in the underlying market.

To justify the necessity of such a distributional assumption on item arrivals, consider the following example in which items arrivals are chosen by an adaptive “adversary” whose goal is to make the buyers’ time-averaged utilities of any online algorithm deviate from the static equilibrium utilities (defined by the “hindsight market,” i.e., the $n \times T$ market with items $\theta_1, \dots, \theta_T$) as much as possible. For simplicity (and w.l.o.g.), budgets, valuations and supplies are not normalized in this example.

- There are n buyers with equal budgets $B_i = 1$ and an item space of $\Theta = [n + 1]$.
- Let the valuation matrix be as follows, for some large $M > n$:

$$v = \begin{bmatrix} 1 & M & M & M & 0 \\ \vdots & \vdots & \vdots & 0 & M \\ \vdots & M & 0 & \vdots & \vdots \\ 1 & 0 & M & M & M \end{bmatrix}.$$

In other words, for item 1, all buyers have valuation 1. For each item $j = 2, \dots, n + 1$, buyer $(n + 2 - j)$ has valuation 0 and other buyers have valuation M .

- The item supplies are 1 for all $j \in [n + 1]$.
- The number of time periods T is large.

Given any fixed $j_0 = 2, \dots, n + 1$, the $n \times T$ static market with $T/2$ items of type 1 and $T/2$ items of type j_0 exhibits the following ME:

- Buyer $i_0 := n + 2 - j_0^*$ receives all $T/2$ items of type 1 with a utility of $u_{i_0}^* = T/2$ and $\beta_{i_0}^* = 2/T$.
- Each buyer $i \in [n] \setminus \{i_0\}$ receives $T/(2(n - 1))$ items of type 2 (i.e., all items of type 2 are evenly distributed among them) with a utility $u_i^* = \frac{MT}{2(n-1)}$ and $\beta_i^* = \frac{2(n-1)}{MT}$.
- The price of item 1 (with $T/2$ copies) is $p_1^* = \max \left\{ 2/T, \frac{2(n-1)}{MT} \right\} = \frac{2}{T}$. The price of item j_0 (with $T/2$ copies) is $p_{j_0}^* = \max \left\{ \beta_{i_0}^* \cdot 0, \frac{2(n-1)}{MT} \cdot M \right\} = 2(n - 1)/T$. To verify these are equilibrium prices, note the following:
 - Buyer i_0 has $v_{i_0 j_0} = 0 < v_{i_0 1} = 1$. Hence, given p^* , will strictly prefer type 1 items over type j_0 . Her equilibrium allocation also consists of only type-1 items that cost exactly her budget: $\frac{T}{2} \cdot \frac{2}{T} = 1$.
 - Buyer $i \neq i_0$ prefers type j_0 over type 1 since the former has a higher value-per-unit-price:

$$\frac{M}{p_{j_0}^*} = \frac{MT}{2(n-1)} > \frac{T}{2} = \frac{1}{p_1^*}.$$

Her equilibrium allocation also consists of only type- j_0 items that cost exactly her budget: $\frac{T}{2(n-1)} \cdot \frac{2(n-1)}{T} = 1$.

Given any online algorithm, consider the following adaptive adversary.

- In the first $T/2$ time steps, every item arrival is of type 1, that is, $\theta_t = 1, t = 1, \dots, T/2$. The algorithm must irrevocably allocate them to the buyers.
- Since every buyer has the same valuation 1 on type 1, there exists a buyer i_0 that receives $u_{i_0} \leq T/(2n)$ up to $T/2$.
- Then, the adversary picks $j_0 = n + 2 - i_0$ and every item arrival in the remaining $T/2$ periods is of type j_0 (which has value zero for i_0).

In this way, buyer i_0 only receives a total utility of u_{i_0} across all T . As shown above, in the static (hindsight) $n \times T$ market, she should have received an equilibrium utility of $u_{i_0}^* = T/2$ (i.e., being allocated all type-1 items). Hence, the realized utility of buyer i_0 is only $1/n$ of her static equilibrium utility.

Reformulation of $(\mathcal{D}_{\text{EG}})$ into (3)

The reformulation is mainly based on [24, §5], except that we now allow non-uniform supplies s instead of $s(\theta) = 1$ for all $\theta \in \Theta$. Assuming uniform supplies is w.l.o.g. in the static Fisher market (via rescaling all v_i) but is not so in OFM, since s in an FOM represents the arbitrary, unknown underlying item distribution from which item arrivals are drawn.

In $(\mathcal{D}_{\text{EG}})$, fixing a $\beta > 0$, setting

$$p = \max_i \beta_i v_i \in L^1(\Theta)_+,$$

i.e., the smallest L^1 function greater than or equal to $\beta_i v_i$ for all i , clearly minimizes the objective subject to the constraints. Hence, we can eliminate p in this way and write $(\mathcal{D}_{\text{EG}})$ as a finite-dimensional convex program in β . Here, $\beta_i v_i, s\rangle$ is convex in β since $\beta \mapsto \max_i \beta_i v_i(\theta)$ is convex

for any $\theta \in \Theta$. More specifically, for any $\beta, \gamma \in \mathbb{R}_+^n$, $\lambda \in [0, 1]$, $\theta \in \Theta$, we have

$$\max_i (\lambda \beta_i + (1 - \lambda) \gamma_i) v_i(\theta) \leq \lambda \max_i \beta_i v_i(\theta) + (1 - \lambda) \max_i \gamma_i v_i(\theta).$$

Hence,

$$\langle \max_i (\lambda \beta + (1 - \lambda) \gamma) v_i, s \rangle \leq \lambda \langle \max_i \beta_i v_i, s \rangle + (1 - \lambda) \langle \max_i \gamma_i v_i, s \rangle.$$

Due to the strong convexity assumption in Theorem 2, we would need the function

$$\beta \mapsto \sum_i B_i \log \beta_i$$

to be strongly convex on its domain. However, it is only strictly but not strongly convex on \mathbb{R}_{++}^n . To resolve this, we use the following lemma. It is similar to [24, Lemma 4] except that we non-uniform supplies.

Lemma 1. *Assume the normalizations in (h) in §2. Then, the equilibrium utilities satisfy $B_i \leq u_i^* \leq 1$ and hence $B_i \leq \beta_i^* = B_i/u_i^* \leq 1$.*

Proof. Since any buyer can get at most the entire set of items (given by the supply s),

$$u_i^* \leq \langle v_i, s \rangle = 1,$$

where the last inequality is due to the normalization $\langle v_i, s \rangle = 1$ in (h) in §2. In any ME (x^*, p^*) , Theorem 1 implies

$$\langle p^*, x_i^* \rangle = \beta_i^* \langle v_i, x_i^* \rangle = \beta_i^* u_i^* = B_i,$$

that is, each buyer i spends her entire budget. Hence, by the normalization $\|B\|_1 = 1$ and market clearance $\langle p^*, s - \sum_i x_i^* \rangle = 0$, we have

$$\langle p^*, s \rangle = \sum_i \langle p^*, x_i^* \rangle = \sum_i B_i = \|B\|_1 = 1 \Rightarrow \langle p^*, B_i s \rangle = B_i.$$

In other words, given item price p^* , each buyer i can afford the proportional allocation $x_i^\circ := B_i s$. Hence, the buyer optimality property of ME implies that buyer i 's equilibrium utility is at least the proportional share:

$$u_i^* \geq \langle v_i, x_i^\circ \rangle = B_i \langle v_i, s \rangle = B_i.$$

Since $B_i \leq u_i^* \leq 1$ and $\beta_i^* = B_i/u_i^*$ at equilibrium (Theorem 1), we have

$$B_i \leq \beta_i^* \leq 1.$$

□

By Lemma 1, adding the constraints

$$B_i/(1 + \delta_0) \leq \beta_i \leq 1 + \delta_0, \forall i$$

to the convex program does not affect its optimal solution β^* . Here, $\delta_0 > 0$ is to ensure $\beta_i^* \in (l_i, h_i)$ (the open interval), which facilitates the convergence analysis of cumulative utilities. To simplify the constants, one can take $\delta_0 = 1$. Numerical experiments suggest that its value does not affect the speeds of convergence of quantities of interest.

Combining the above yields the reformulation (3). To align (3) with (2), for each $\theta \in \Theta$ (corresponding to Z in §4) and $\beta \in \mathbb{R}_+^n$ (corresponding to w in §4), let

$$f_\theta(\beta) := \max_i \beta_i v_{ij}.$$

Then,

$$f(\beta) := \mathbf{E} f_\theta(\beta) = \langle \max_i \beta_i v_i, s \rangle,$$

where the expectation is over $\theta \sim s$, i.e., a random variable with distribution s (corresponding to $z \sim \mathcal{D}$ in §4).

Proof of Theorem 3

It follows immediately from Theorem 2, as long as the function

$$\Psi(\beta) = - \sum_i B_i \log \beta_i$$

is strongly convex modulo σ and $\mathbf{E}\|v_{i_t}(\theta_t)\mathbf{e}^{(i_t)}\|^2 \leq G^2$. We now show them. Note that Ψ is twice differentiable and has a diagonal Hessian

$$\nabla^2 \Psi(\beta) = \begin{bmatrix} \frac{B_1}{\beta_1^2} & & \\ & \ddots & \\ & & \frac{B_n}{\beta_n^2} \end{bmatrix}$$

at any $\beta > 0$. Clearly, its smallest eigenvalue can be bounded as

$$\lambda_{\min}(\nabla^2 \Psi(\beta)) \geq \min_i \frac{B_i}{\beta_i}.$$

Denote $\kappa = 1/(\min_i B_i)$. For any β feasible to (3), by the constraints $B_i/(1 + \delta_0) \leq \beta_i \leq 1 + \delta_0$, we have

$$\lambda_{\min}(\nabla^2 \Psi(\beta)) \geq \min_i \min_{\beta_i \in [B_i/(1+\delta_0), 1+\delta_0]} \frac{B_i}{\beta_i^2} = \min_i \frac{B_i}{(1 + \delta_0)^2} = \frac{1}{\kappa(1 + \delta_0)^2}.$$

Therefore, Ψ is strongly convex on $[B/(1 + \delta_0), (1 + \delta_0)\mathbf{1}]$ with modulus $\sigma = \frac{1}{\kappa(1+\delta_0)^2}$. Finally, we have

$$\mathbf{E}\|v_{i_t}\mathbf{e}^{(i_t)}\|^2 \leq \max_i \mathbf{E}_{\theta \sim s}[v_i(\theta)^2] = G^2 \leq \max_i \|v_i\|_\infty^2.$$

Proof of Theorem 4

Intuitively, our proof uses the fact that if β_i^t and β_i^* are near each other, then $\frac{B_i}{\beta_i^t}$ will be near $\frac{B_i}{\beta_i^*} = u_i^*$ as well. Recall that $g_i^t = u_i^t$ (i.e., the subgradient of $\beta \mapsto \max_i \beta_i v_i(\theta_t)$ that we choose corresponds to the utility buyer i receives at time t) and hence $\bar{g}_i^t = u_i^t$. Since

$$\beta^{t+1} = \Pi_{[l_i, h_i]} \left(\frac{B_i}{\bar{g}_i^t} \right),$$

we know that if no projection occurs (i.e., if $\frac{B_i}{\bar{g}_i^t} \in [l_i, h_i]$) at iteration t , then

$$\frac{B_i}{\beta_i^{t+1}} = \bar{g}_i^t.$$

Thus, we split our proof into two cases: the case where projection occurs (i.e., $\frac{B_i}{\bar{g}_i^t} \notin [l_i, h_i]$), and the case where projection does not occur. As we will see, the probability of a projection at time step t converges to 0 as t grows.

For each i , consider the event that no projection occurs:

$$A_i^t := \{l_i \leq B_i/\bar{g}_i^t \leq h_i\}.$$

Conditioning on the complementary event $(A_i^t)^c = \{\bar{g}_i^t \notin [l_i, h_i]\}$, it holds that

$$|\beta_i^{t+1} - \beta_i^*| > \epsilon_i \Rightarrow \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2 \geq \mathbf{P}[(A_i^t)^c] \epsilon_i^2 \Rightarrow \mathbf{P}[(A_i^t)^c] \leq \frac{1}{\epsilon_i^2} \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2.$$

Conditioning on A_i^t , we have $B_i/\bar{g}_i^t = \beta_i^{t+1}$. Furthermore, since

$$0 \leq \bar{g}_i^t = \frac{1}{t} \sum_{\tau=1}^t v_{ij_\tau} \mathbb{I}\{i = i_\tau\} \leq \|v_i\|_\infty$$

and $\|v_i\|_\infty \geq 1 \geq u_i^*$, we have the following upper bound on the difference between the time average of realized utilities and the equilibrium utility of buyer i :

$$|\bar{g}_i^t - u_i^*| \leq \max\{u_i^*, \|v_i\|_\infty\} = \|v_i\|_\infty.$$

Now, splitting the expectation by the two complementary events A_i^t and $(A_i^t)^c$, we can apply the above bounds to get

$$\begin{aligned} \mathbf{E}(\bar{g}_i^t - u_i^*)^2 &= \mathbf{E}[\mathbb{I}_{(A_i^t)^c} \cdot (\bar{g}_i^t - u_i^*)^2] + \mathbf{E}\left[\mathbb{I}_{A_i^t} \cdot \left(\frac{B_i}{\beta_i^{t+1}} - u_i^*\right)^2\right] \\ &\leq \|v_i\|_\infty^2 \mathbf{E}[\mathbb{I}_{(A_i^t)^c}] + (u_i^*)^2 \mathbf{E}\left[\mathbb{I}_{A_i^t} \cdot \left(\frac{B_i}{\beta_i^{t+1} u_i^*} - 1\right)^2\right] \\ &\leq \|v_i\|_\infty^2 \mathbf{P}[(A_i^t)^c] + (u_i^*)^2 \cdot \mathbf{E}\left(\frac{\beta_i^{t+1} - \beta_i^*}{\beta_i^{t+1}}\right)^2 \\ &\leq \frac{\|v_i\|_\infty^2}{\epsilon_i^2} \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2 + \left(\frac{(1 + \delta_0)u_i^*}{B_i}\right)^2 \cdot \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2 \\ &\leq \left(\frac{\|v_i\|_\infty^2}{\epsilon_i^2} + \left(\frac{1 + \delta_0}{B_i}\right)^2\right) \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2. \end{aligned}$$

Since $B_i \leq \beta_i^* \leq 1$, we have $(\kappa := 1/(\min_i B_i))$

$$\epsilon_i \geq B_i \delta_0 / (1 + \delta_0) > \delta_0 / \kappa > 0.$$

Summing up across all i , using Theorem 3 and the above bound, we get

$$\begin{aligned} \mathbf{E}\|\bar{g}^t - u^*\|^2 &\leq \sum_i \left(\frac{\|v_i\|_\infty^2}{\epsilon_i^2} + \left(\frac{1 + \delta_0}{B_i}\right)^2 \right) \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2 \\ &\leq \left(\|v\|_\infty^2 \left(\frac{\kappa}{\delta_0}\right)^2 + ((1 + \delta_0)\kappa)^2 \right) \sum_i \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2 \\ &\leq \left(\|v\|_\infty^2 \left(\frac{\kappa}{\delta_0}\right)^2 + ((1 + \delta_0)\kappa)^2 \right) \frac{(6 + \log(t+1))G^2}{(t+1)\sigma^2} \\ &= C \cdot \frac{(6 + \log(t+1))G^2}{(t+1)\sigma^2}. \end{aligned}$$

Proof of Theorem 5

First, note that \bar{b}_i^t can be decomposed as follows.

$$\begin{aligned} \bar{b}_i^t &= \frac{1}{t} \sum_{\tau=1}^t \beta_i^\tau v_{ij_\tau} \mathbb{I}\{i = i_\tau\} \\ &= \beta_i^* \cdot \frac{1}{t} \sum_{\tau=1}^t v_{ij_\tau} \mathbb{I}\{i = i_\tau\} + \frac{1}{t} \sum_{\tau=1}^t (\beta_i^\tau - \beta_i^*) v_{ij_\tau} \mathbb{I}\{i = i_\tau\} \\ &= \beta_i^* \bar{g}_i^t + \frac{1}{t} \sum_{\tau=1}^t (\beta_i^\tau - \beta_i^*) v_{ij_\tau} \mathbb{I}\{i = i_\tau\}. \end{aligned}$$

Next, we bound the second term as follows, using convexity of $(\cdot)^2$ and $\|v_{ij_\tau}\| \leq \|v_i\|_\infty$:

$$\left(\frac{1}{t} \sum_{\tau=1}^t (\beta_i^\tau - \beta_i^*) v_{ij_\tau} \mathbb{I}\{i = i_\tau\} \right)^2 \leq \frac{1}{t} \sum_{\tau=1}^t (\beta_i^\tau - \beta_i^*)^2 \|v_i\|_\infty^2.$$

Then, we bound the square difference between expenditure and budget as follows, using $(x + y)^2 \leq 2(x^2 + y^2)$ for any $x, y \in \mathbb{R}$:

$$(\bar{b}_i^t - B_i)^2 \leq 2 \left[(\beta_i^* \bar{g}_i^t - B_i)^2 + \left(\frac{1}{t} \sum_{\tau=1}^t (\beta_i^\tau - \beta_i^*) v_{ij_\tau} \mathbb{I}\{i = i_\tau\} \right)^2 \right].$$

Combining the above two inequalities, taking expectation on both sides and using $\beta_i^* = B_i/u_i^*$, we have

$$\mathbf{E}(\bar{b}_i^t - B_i)^2 \leq 2 \left[(\beta_i^*)^2 \mathbf{E}(\bar{g}_i^t - u_i^*)^2 + \|v_i\|_\infty^2 \frac{1}{t} \sum_{\tau=1}^t \mathbf{E}(\beta_i^\tau - \beta_i^*)^2 \right]. \quad (5)$$

When $t \geq 3$, we have $\frac{\log(t+1)}{t+1} < \frac{\log t}{t}$ (since $(\frac{\log t}{t})' = \frac{1-\log t}{t^2} < 0$ for all $t \geq 3$). By the proof of [48, Corollary 4],

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

Finally, summing up (5) across all i , using $\beta_i^* \leq 1$, Theorems 3 and 4, and (6), we have

$$\begin{aligned} \mathbf{E} \|\bar{b}^t - B\|^2 &\leq 2 \left[\mathbf{E} \|\bar{g}^t - u^*\|^2 + \|v\|_\infty^2 \frac{1}{t} \sum_{\tau=1}^t \mathbf{E} \|\beta^\tau - \beta^*\|^2 \right] \\ &\leq 2 \left[C \cdot \frac{(6 + \log t) G^2}{t \sigma^2} + \|v\|_\infty^2 \frac{1}{t} \left(6(1 + \log t) + \frac{(\log t)^2}{2} \right) \frac{G^2}{\sigma^2} \right] \\ &= \frac{2G^2}{t \sigma^2} \left(6(C + \|v\|_\infty^2) + (C + 6\|v\|_\infty^2) \log t + \frac{\|v\|_\infty^2}{2} (\log t)^2 \right). \end{aligned}$$

Proof of Theorem 6

For any $\theta \in \Theta$, since $\|\cdot\|_\infty$ is 1-Lipschitz continuous w.r.t. itself, we have

$$\begin{aligned} \left| p^*(\theta) - \max_i \beta_i^t v_i(\theta) \right| &\leq \left| \max_i \beta_i^* v_i(\theta) - \max_i \beta_i^t v_i(\theta) \right| \\ &\leq \max_i |\beta_i^* v_i(\theta) - \beta_i^t v_i(\theta)| \\ &\leq \|v\|_\infty \|\beta^t - \beta^*\|_\infty. \end{aligned} \quad (7)$$

Analysis of regret r_i^t . Let $(z_i^\tau)_{\tau \in [t]} \in [0, 1]^t$ be any feasible allocation on the arrived items θ_τ , $\tau \in [t]$ such that

$$\frac{1}{t} \sum_{\tau=1}^t p^\tau(\theta_t) z_i^\tau \leq B_i.$$

Using $p^\tau(\theta_\tau) = \max_i \beta_i^\tau v_i(\theta_\tau)$, we have

$$\begin{aligned} \frac{1}{t} \sum_{\tau=1}^t p^*(\theta_\tau) z_i^\tau &= \frac{1}{t} \sum_{\tau=1}^t p^\tau(\theta_\tau) z_i^\tau + \frac{1}{t} \sum_{\tau=1}^t (p^*(\theta_\tau) - p^\tau(\theta_\tau)) z_i^\tau \\ &\leq B_i + \frac{1}{t} \|v\|_\infty \sum_{\tau=1}^t \|\beta^\tau - \beta^*\|_\infty \quad [\text{by (7) and } 0 \leq z_{ij}^\tau \leq 1]. \end{aligned} \quad (8)$$

Denote

$$\gamma_t = \frac{1}{t} \|v\|_\infty \sum_{\tau=1}^t \|\beta^\tau - \beta^*\|_\infty.$$

In a static ME, by Theorem 1 and the constraints in $(\mathcal{D}_{\text{EG}})$, we have $p^* \geq \beta_i^* v_i$. Hence,

$$\frac{1}{t} \sum_{\tau=1}^t p^*(\theta_\tau) z_i^\tau \geq \beta_i^* \left(\frac{1}{t} \sum_{\tau=1}^t v_i(\theta_\tau) z_i^\tau \right). \quad (9)$$

By (8), (9), $u_i^* = B_i/\beta_i^*$ (Theorem 1) and the definition of ξ_i^t , we have

$$\frac{1}{t} \sum_{\tau=1}^t v_i(\theta_\tau) z_i^\tau \leq \frac{1}{\beta_i^*} (B_i + \gamma_t) = u_i^* \left(1 + \frac{\gamma_t}{B_i}\right) \leq u_i^* + \frac{\gamma_t}{B_i} \leq \bar{u}_i^t + \xi_i^t + \frac{\gamma_t}{B_i}.$$

Hence, the utility level \hat{U}_i^t (Definition 3) satisfies

$$\hat{U}_i^t \leq \bar{u}_i^t + \xi_i^t + \frac{\gamma_t}{B_i}. \quad (10)$$

Note that $\mathbf{E}(\gamma_t^2)$ can be bounded as follows:

$$\mathbf{E}(\gamma_t^2) \leq \|v\|_\infty^2 \frac{1}{t} \sum_{\tau=1}^t \mathbf{E} \|\beta^\tau - \beta^*\|^2 \leq \frac{\|v\|_\infty^2}{t} \left(6(1 + \log t) + \frac{(\log t)^2}{2}\right) \frac{G^2}{\sigma^2} = O\left(\frac{(\log t)^2}{t}\right). \quad (11)$$

where the second inequality is due to Theorem 3 and (6). Combining (10), (11) and $\mathbf{E}(\xi_i^t)^2 = O((\log t)/t)$ (Theorem 4), we have

$$\mathbf{E}(r_i^t)^2 \leq 2 \left(\mathbf{E}(\xi_i^t)^2 + \frac{1}{B_i^2} \mathbf{E}(\gamma_t^2) \right) = O\left(\frac{(\log t)^2}{t}\right).$$

Analysis of envy ρ_i^t . Let $p^* = \max_i \beta_i^* v_i$ (a.e.) be the equilibrium prices. Similar to 8, for any i , we have

$$\begin{aligned} \frac{1}{t} \sum_{\tau=1}^t p^*(\theta_\tau) x_i^\tau &= \bar{b}_i^t + \frac{1}{t} \sum_{\tau=1}^t (p^*(\theta_\tau) - \beta_i^\tau v_i(\theta_\tau)) x_i^\tau \\ &\leq B_i + \Delta_i^t + \eta_i^t. \end{aligned} \quad (12)$$

Using the above (replacing i with k) and $p^* \geq \beta_i^* v_i$, we have

$$\beta_i^* \bar{u}_{ik}^t = \frac{1}{t} \sum_{\tau=1}^t \beta_i^* v_i(\theta_\tau) x_k^\tau \leq \frac{1}{t} \sum_{\tau=1}^t p^*(\theta_\tau) x_k^\tau \leq B_k + \Delta_k^t + \eta_k^t.$$

Hence, using $u_i^* = B_i/\beta_i^* \leq 1$ (Theorem 1 and Lemma 1),

$$\begin{aligned} \frac{\bar{u}_{ik}^t}{B_k} &\leq \frac{1}{B_k} \cdot \frac{1}{\beta_i^*} (B_k + \Delta_k^t + \eta_k^t) \\ &\leq \frac{u_i^*}{B_i} \left(1 + \frac{\Delta_k^t + \eta_k^t}{B_k}\right) \\ &\leq \frac{\bar{u}_i^t}{B_i} + \frac{\xi_i^t}{B_i} + \frac{\Delta_k^t + \eta_k^t}{B_k B_i} \quad [\text{by definition of } \xi_i^t]. \end{aligned}$$

Using the above inequality, we can bound the envy as follows:

$$\rho_i^t \leq \frac{\xi_i^t}{B_i} + \frac{1}{B_i} \max_k \frac{\Delta_k^t + \eta_k^t}{B_k} \leq \kappa \xi_i^t + \kappa^2 \max_k (\Delta_k^t + \eta_k^t). \quad (13)$$

Next, we show the convergence of η_i^t . By (7), we have

$$|\eta_i^t| \leq \sum_{\ell} |\eta_\ell^t| \leq \frac{1}{t} \sum_{\tau=1}^t |p_{j_\tau}^* - \beta_{i_\tau}^\tau v_{i_\tau j_\tau}| \leq \frac{1}{t} \sum_{\tau=1}^t \|v\|_\infty \|\beta^\tau - \beta^*\|_\infty = \gamma_t. \quad (14)$$

Hence, same as (11),

$$\mathbf{E}(\eta_i^t)^2 \leq \|v\|_\infty^2 \frac{1}{t} \sum_{\tau=1}^t \mathbf{E} \|\beta^\tau - \beta^*\|^2 \leq \|v\|_\infty^2 \frac{1}{t} \left(6(1 + \log t) + \frac{(\log t)^2}{2}\right) \frac{G^2}{\sigma^2}. \quad (15)$$

By Theorems 3 and 5, we know that $\mathbf{E}(\xi_i^t)^2 = O((\log t)/t)$ and $\mathbf{E}(\Delta_i^t)^2 = O((\log t)^2/t)$. Together with (15) and (13), we have

$$\mathbf{E}(\rho_i^t)^2 \leq \kappa \mathbf{E}(\xi_i^t)^2 + \kappa^2 \sum_{\ell} (\mathbf{E}(\Delta_\ell^t)^2 + \mathbf{E}(\eta_\ell^t)^2) = O\left(\frac{(\log t)^2}{t}\right).$$

C Extension to Quasilinear Utilities

We show that PACE can be easily extended to the case of a *quasilinear* (QL) market (i.e., where buyers have QL utilities). We show that most of the convergence results in §5 still holds. The static quasilinear market setup is the same as the linear case in §2 (which allows a possibly infinite item space Θ), except the following:

- For given item prices $p \in L_+^1$, each buyer i has a quasilinear utility function, i.e.,

$$u_i(x_i) = \langle v_i, x_i \rangle - \langle p, x_i \rangle.$$

- Without loss of generality, assume $\|B\|_1 = 1$ and all buyers' valuations are nontrivial, i.e., $\langle v_i, s \rangle > 0$ for all i . Due to the structure of QL utilities, we cannot normalize the valuations v_i and budgets B_i separately without loss of generality. Instead, they can only be scaled at the same time by the same constant.

Same as before, each buyer i has a budget $B_i > 0$ and can only choose among budget-feasible allocations, that is, x_i such that $\langle p, x_i \rangle \leq B_i$. In this case, an allocation-price pair (x^*, p^*) is a *quasilinear market equilibrium* (QLME) if the following holds (see [24, §6] and [17, §4]):

- Buyers are optimal: $x_i^* \in D_i(p^*) := \arg \max\{\langle v_i - p^*, x_i \rangle : x_i \in L_+^\infty, \langle p^*, x_i \rangle \leq B_i\}$.
- The market clears: $\sum_i x_i \leq s$ and $\langle p^*, s - \sum_i x_i^* \rangle = 0$.

As shown in [24, §6], the following pair of (possibly infinite-dimensional) convex programs capture QLME:⁷

$$\begin{aligned} \sup \quad & \sum_i (B_i \log u_i - \delta_i) \\ \text{s.t.} \quad & u_i \leq \langle v_i, x_i \rangle + \delta_i, \forall i \in [n], \\ & \sum_i x_i \leq s, \\ & u_i \geq 0, \delta_i \geq 0, x_i \in L_1(\Theta)_+, \forall i \in [n]. \end{aligned} \tag{P_{QLME}}$$

$$\begin{aligned} \inf \quad & \langle p, s \rangle - \sum_i B_i \log \beta_i \\ \text{s.t.} \quad & p \geq \beta_i v_i, \beta_i \leq 1, \forall i \in [n], \\ & p \in L_1(\Theta)_+, \beta \in \mathbb{R}_+^d. \end{aligned} \tag{D_{QLME}}$$

In the sequel, we use (x^*, u^*, δ^*) to denote an optimal solution of (P_{QLME}) (in which u^* and δ^* are unique) and (p^*, β^*) to denote the optimal solution of (D_{QLME}) . As shown in [24, §6], the following KKT conditions of (P_{QLME}) and (D_{QLME}) are necessary and sufficient for (x^*, p^*) being a QLME.

- $\delta_i^*(1 - \beta_i^*) = 0$ for all i (complementary slackness).
- $u_i^* = B_i / \beta_i^*$ for all i .
- $p^* = \max_i \beta_i^* v_i$ (a.e.) for all j .
- $\langle p^* - \beta_i^* v_i, x_i^* \rangle = 0$ for all i .

Let (x^*, p^*) denote a QLME. The equilibrium utility of buyer i (i.e., the amount of utility buyer i receives at a QLME) is

$$u_i^{\text{QLME}} := \langle v_i - p^*, x_i^* \rangle = (1 - \beta_i^*) \langle v_i, x_i^* \rangle = (1 - \beta_i^*) (u_i^* - \delta_i^*),$$

which is unique and does not depend on the choice of the equilibrium allocation x^* . In general, u_i^{QLME} is not the same as u_i^* in the optimal solution of (P_{QLME}) . In comparison, the term $\langle v_i, x_i^* \rangle$

⁷There, the authors assume $s = \mathbf{1}$, which is w.l.o.g. for static Fisher markets. Similar to the case of Theorem 1 for linear utilities, all results can be easily extended to the case of $s \in L_+^\infty$.

can be viewed as the equilibrium gross utility before subtracting the price $\langle p^*, x_i^* \rangle$ of the allocation x_i^* . The above equilibrium quantities satisfy the following [24, §6].

- If $\beta_i^* = 1$, then $\langle p^* - \beta_i^* v_i, x_i^* \rangle = 0$ implies its gross utility and expenditure are equal, which give an equilibrium utility of zero:

$$\langle v_i, x_i^* \rangle = \beta_i^* \langle v_i, x_i^* \rangle = \langle p^*, x_i^* \rangle = u_i^* - \delta_i^* \Rightarrow u_i^{\text{QLME}} = \langle v_i - p^*, x_i^* \rangle = 0.$$

- If $\beta_i^* < 1$, then $\delta_i^* = 0$ by complementary slackness (the first KKT condition above). Hence, the gross utility is $\langle v_i, x_i^* \rangle = u_i^*$ and

$$u_i^{\text{QLME}} = \langle v_i - p^*, x_i^* \rangle = (1 - \beta_i^*) \langle v_i, x_i^* \rangle = (1 - \beta_i^*) u_i^*.$$

Similar to the proof of [23, Lemma 5], we can show that

$$u_i^* \leq \langle v_i, s \rangle + B_i.$$

Hence,

$$\beta^* = \frac{B_i}{u_i^*} \geq \frac{B_i}{\langle v_i, s \rangle + B_i} > \beta_i^{\min} := \frac{B_i}{\langle v_i, s \rangle + 2B_i} > 0.$$

The choice of β_i^{\min} is to ensure that $\beta_i^* - \beta_i^{\min} > 0$, which simplifies the analysis of the dynamics. Substituting $p = \max_i \beta_i v_i$ and using the bounds $\beta_i^{\min} \leq \beta_i^* \leq 1$, we can solve the following convex program for the equilibrium utility prices β^* , where $\beta^{\min} := (\beta_1^{\min}, \dots, \beta_n^{\min})$:

$$\min_{\beta \in [\beta^{\min}, 1]} \langle p, s \rangle - \sum_i B_i \log \beta_i. \quad (16)$$

Applying dual averaging to the convex program (16), we arrive at the following PACE dynamics for an online QL market (i.e., an OFM with buyers having QL utilities). Set the initial pacing multipliers are $\beta_i^1 = 1$ for all i . At time t , the following steps take place.

- An item $\theta_t \in \Theta$ arrives, which determines a winner $i_t = \min \arg \max_i \beta_i^t v_i(\theta_t)$.
- The stochastic subgradient is $g^t = v_{i_t}(\theta_t) \mathbf{e}^{(i_t)}$, or $g_i^t = v_i(\theta_t) \mathbb{I}\{i = i_t\}$ for each i .
- Each buyer i pays a price (expenditure)

$$b_i^t = \beta_i^t v_i(\theta_t) \mathbb{I}\{i = i_t\}$$

and receives a (net) utility of

$$u_i^t = g_i^t - b_i^t = (1 - \beta_i^t) v_i(\theta_t) \mathbb{I}\{i = i_t\}, \quad (17)$$

which is the value of the item minus the price paid. Here, only the winning buyer i_t may get a potentially nonzero utility $u_{i_t}^t$; other buyers $i \neq i_t$ gets 0 (and pays zero).

- Update the dual average: for each i , $\bar{g}^t = \frac{t-1}{t} \bar{g}^{t-1} + \frac{1}{t} g^t$, which ensures $\bar{g}_i^t = \frac{1}{t} \sum_{\tau=1}^t v_{ij_\tau} \mathbb{I}\{i = i_\tau\}$ for all i (same as in the linear case).
- Compute the next pacing multiplier (similar to the linear case, except the lower bound for β_i^t being β_i^{\min} instead of B_i):

$$\beta_i^{t+1} = \arg \min_{\beta_i \in [\beta_i^{\min}, 1]} \{\bar{g}_i^t \beta_i - B_i \log \beta_i\} \Rightarrow \beta_i^{t+1} = \Pi_{[\beta_i^{\min}, 1]} \left(\frac{B_i}{\bar{g}_i^t} \right).$$

Same as in the linear case, we do not need any distributional assumption on the item arrivals to run PACE. In subsequent convergence analysis, however, we assume that the items θ_t are drawn i.i.d. from a distribution s (i.e., $s \in L_+^\infty$ and $s(\Theta) = 1$). We also assume that $v_i \in L_+^\infty$ for all i . Let (x^*, p^*) denote a QLME of the underlying static QL market with supplies s .

Convergence of QL pacing multipliers. Analogous to Theorem 3, in the QL case, we can show that the pacing multipliers β^t converge to the equilibrium utility prices β^* in mean square. It is a direct consequence of the general convergence result of dual averaging (Theorem 2).

Theorem 7. For $t = 1, 2, \dots$, it holds that

$$\mathbf{E}\|\beta^t - \beta^*\|^2 \leq \frac{(6 + \log t)G^2}{t\sigma^2},$$

where G^2, κ are the same as in Theorem 3 and

$$\sigma = \min_i \min_{\beta_i \in [\beta_i^{\min}, 1]} \frac{B_i}{\beta_i^2} = \min_i B_i = \frac{1}{\kappa}.$$

Convergence of QL utilities and expenditures. Next, we show mean-square convergence of time-averaged utilities $\bar{u}^t = \frac{1}{t} \sum_{\tau=1}^t u_i^\tau$ and expenditures $\bar{b}^t = \frac{1}{t} \sum_{\tau=1}^t \beta_i^\tau v_{ij_\tau} \mathbb{I}\{i = i_\tau\}$. In the QL case, the dual average g_i^t can also be viewed as the per-period gross utility before subtracting the price $\beta_i^t v_i(\theta_t)$. In this way, \bar{g}_i^t is the time-averaged gross utility of buyer i .

Theorem 8. For $t = 1, 2, \dots$ and each i , the following holds.

- If $\beta_i^* = 1$, then $u_i^{\text{QLME}} = 0$ and $B_i = u_i^*$. In this case,

$$\mathbf{E}\left(u_i^t - u_i^{\text{QLME}}\right)^2 = \mathbf{E}|u_i^t|^2 \leq \|v_i\|_\infty^2 \mathbf{E}(1 - \beta_i^t)^2 = O\left(\frac{\log t}{t}\right). \quad (18)$$

- If $\beta_i^* < 1$, then $u_i^{\text{QLME}} > 0$, $\delta_i^* = 0$ and $u^{\text{QLME}} = (1 - \beta_i^*)u_i^*$. In this case, the gross utility \bar{g}_i^t , realized (net) utility \bar{u}_i^t and expenditures \bar{b}_i^t converge as follows:

$$\begin{aligned} \mathbf{E}(\bar{g}_i^t - u_i^*)^2 &\leq C_i \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2 = O\left(\frac{\log t}{t}\right), \\ \mathbf{E}\left(\bar{u}_i^t - u_i^{\text{QLME}}\right)^2 &\leq R_i^t, \quad \mathbf{E}(\bar{b}_i^t - B_i)^2 \leq R_i^t, \end{aligned}$$

where

$$\begin{aligned} C_i &= \frac{\|v_i\|_\infty^2}{\epsilon_i^2} + \frac{\left(\frac{\|v_i\|_1}{m} + 2B_i\right)^4}{B_i^2}, \quad \epsilon_i = \min\{1 - \beta_i^*, \beta_i^* - \beta_i^{\min}\}, \\ R_i^t &= 2 \left[\mathbf{E}(\bar{g}_i^t - u_i^*)^2 + \frac{\|v_i\|_\infty^2}{t} \sum_{\tau=1}^t \mathbf{E}(\beta_i^\tau - \beta_i^*)^2 \right], \quad \mathbf{E}R_i^t = O\left(\frac{(\log t)^2}{t}\right). \end{aligned}$$

Hence, the mean-square error $\mathbf{E}\|\bar{u}^t - u^{\text{QLME}}\|^2$ is either $O((\log t)^2/t)$ (when some $\beta_i^* < 1$) or $O((\log t)/t)$ (when all $\beta_i^* = 1$).

Proof. The $\beta_i^* = 1$ case is clear. We prove the $\beta_i^* < 1$ case.

Convergence of \bar{g}_i^t . Denote the event $A_i^t = \{B_i \leq \bar{g}_i^t \leq B_i/\beta_i^{\min}\}$. Then, $(A_i^t)^c$ means $B_i/\bar{g}_i^t \notin [\beta_i^{\min}, 1]$ and hence $|\beta_i^{t+1} - \beta_i^*| > \epsilon_i$. Similar to the linear case, we deduce

$$\mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2 \geq \mathbf{P}[(A_i^t)^c] \epsilon_i^2 \Rightarrow \mathbf{P}[(A_i^t)^c] \leq \frac{1}{\epsilon_i^2} \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2.$$

Furthermore, since $0 \leq \bar{g}_i^t \leq \|v_i\|_\infty$ (same as in the linear case) and $u_i^* \leq \langle v_i, s \rangle + B_i$, we have

$$\begin{aligned}
\mathbf{E}(\bar{g}_i^t - u_i^*)^2 &= \mathbf{E}[\mathbb{I}_{(A_i^t)^c} \cdot (\bar{g}_i^t - u_i^*)^2] + \mathbf{E}\left[\mathbb{I}_{A_i^t} \cdot \left(\frac{B_i}{\beta_i^{t+1}} - u_i^*\right)^2\right] \\
&\leq \|v_i\|_\infty^2 \mathbf{E}[\mathbb{I}_{(A_i^t)^c}] + (u_i^*)^2 \mathbf{E}\left[\mathbb{I}_{A_i^t} \cdot \left(\frac{B_i}{\beta_i^{t+1} u_i^*} - 1\right)^2\right] \\
&\leq \|v_i\|_\infty^2 \mathbf{P}[(A_i^t)^c] + (u_i^*)^2 \cdot \mathbf{E}\left(\frac{\beta_i^{t+1} - \beta_i^*}{\beta_i^{t+1}}\right)^2 \\
&\leq \frac{\|v_i\|_\infty^2}{\epsilon_i^2} \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2 + \left(\frac{u_i^*}{\beta_i^{\min}}\right)^2 \cdot \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2 \\
&\leq \left(\frac{\|v_i\|_\infty^2}{\epsilon_i^2} + \left(\frac{(\langle v_i, s \rangle + B_i)(\langle v_i, s \rangle + 2B_i)}{B_i}\right)^2\right) \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2 \\
&= C_i \mathbf{E}(\beta_i^{t+1} - \beta_i^*)^2 = O\left(\frac{\log t}{t}\right).
\end{aligned}$$

Convergence of expenditures \bar{b}_i^t . Similar to the linear case, note that \bar{b}_i^t can be decomposed as follows (where $x_i^\tau := \mathbb{I}\{i = i_\tau\}$ denotes whether buyer i wins at time step τ):

$$\bar{b}_i^t := \frac{1}{t} \sum_{\tau=1}^t \beta_i^\tau v_i(\theta_\tau) x_i^\tau = \beta_i^* \bar{g}_i^t + \frac{1}{t} \sum_{\tau=1}^t (\beta_i^\tau - \beta_i^*) v_i(\theta_\tau) x_i^\tau.$$

Hence, using $\beta_i^* = B_i/u_i^* \leq 1$, $(x+y)^2 \leq 2(x^2+y^2)$, convexity of $(\cdot)^2$ and $v_i(\theta_\tau) x_i^\tau \leq \|v_i\|_\infty$, we have

$$\begin{aligned}
\mathbf{E}(\bar{b}_i^t - B_i)^2 &\leq 2 \left[\mathbf{E}(\beta_i^* \bar{g}_i^t - B_i)^2 + \mathbf{E}\left(\frac{1}{t} \sum_{\tau=1}^t (\beta_i^\tau - \beta_i^*) v_i(\theta_\tau) x_i^\tau\right)^2 \right] \\
&\leq 2 \left[(\beta_i^*)^2 \mathbf{E}(\bar{g}_i^t - u_i^*)^2 + \frac{\|v_i\|_\infty^2}{t} \sum_{\tau=1}^t \mathbf{E}(\beta_i^\tau - \beta_i^*)^2 \right] \\
&\leq 2 \left[\mathbf{E}(\bar{g}_i^t - u_i^*)^2 + \frac{\|v_i\|_\infty^2}{t} \sum_{\tau=1}^t \mathbf{E}(\beta_i^\tau - \beta_i^*)^2 \right] = R_i^t. \tag{19}
\end{aligned}$$

The order of $\mathbf{E}R_i^t$ is given by those of $\mathbf{E}(\bar{g}_i^t - u_i^*)^2$ and $\sum_{\tau=1}^t \mathbf{E}(\beta_i^\tau - \beta_i^*)^2$, which are $O((\log t)/t)$ and $O((\log t)^2/t)$, respectively.

Convergence of utilities \bar{u}_i^t . For a buyer i with $\beta_i^* < 1$, let

$$\epsilon_i = \min\{1 - \beta_i^*, \beta_i^* - \beta_i^{\min}\} > 0.^8$$

Express \bar{u}_i^t as follows:

$$\begin{aligned}
\bar{u}_i^t &= \frac{1}{t} \sum_{\tau=1}^t (1 - \beta_i^\tau) v_i(\theta_\tau) x_i^\tau \\
&= (1 - \beta_i^*) \bar{g}_i^t + \frac{1}{t} \sum_{\tau=1}^t (\beta_i^\tau - \beta_i^*) v_i(\theta_\tau) x_i^\tau.
\end{aligned}$$

⁸The analysis of this case also works for $\beta_i^* = 0$ but its resulting bound is not as tight as the above one for the case $\beta_i^* = 1$.

Since $u_i^{\text{QLME}} = (1 - \beta_i^*)u_i^*$, similar to (19), we have

$$\begin{aligned} \mathbf{E} \left(\bar{u}_i^t - u_i^{\text{QLME}} \right)^2 &\leq 2 \left[(1 - \beta_i^*) \mathbf{E}(\bar{g}_i^t - u_i^*)^2 + \mathbf{E} \left(\frac{1}{t} \sum_{\tau=1}^t (\beta_i^\tau - \beta_i^*) v_i(\theta_\tau) x_i^\tau \right)^2 \right] \\ &\leq 2 \left[\mathbf{E}(\bar{g}_i^t - u_i^*)^2 + \frac{\|v_i\|_\infty^2}{t} \sum_{\tau=1}^t \mathbf{E}(\beta_i^\tau - \beta_i^*)^2 \right] = R_i^t. \end{aligned}$$

Finally, if all $\beta_i^* = 1$, (18) implies that $\mathbf{E} \|\bar{u}^t - \bar{u}^{\text{QLME}}\|^2 = O((\log t)/t)$. If some $\beta_i^* < 1$, since $\mathbf{E} R_i^t = O((\log t)^2/t)$, so is $\mathbf{E} \|\bar{u}^t - \bar{u}^{\text{QLME}}\|^2$. \square

D More Details on the Experiments

In each experiment, we will have some underlying valuations, items will be drawn one-at-a-time, uniformly at random, from the set of possible items, on which we run the PACE dynamics. We have several outcome measures of interest for asking how close we are to the static equilibrium quantities at each point. First, we look at convergence of realized utilities. In each case we consider the realized utilities up to time t and look at the deviation from equilibrium utility normalized by the equilibrium utility level. We look at both the average and the worst-case deviations. Formally these are calculated as $\|(\bar{u}^t - u^*)/u^*\|_1/n$ for the average deviation and $\|(\bar{u}^t - u^*)/u^*\|_\infty$ for the maximum (over buyers) deviation. We also measure deviations of the pacing multiplier β^t from β^* and deviations of time-averaged cumulative expenditure \bar{b}^t from buyers' budgets $B = (B_1, \dots, B_n)$ using analogous normalizations. In the plots, we add horizontal lines for the same error measures for the proportional shares of the static underlying Fisher market (each buyer receiving B_i of each item), a 'baseline' solution.

We consider 3 different market datasets. The first two datasets are recommender systems which we turn into markets. The final is taken from a survey experiment. We point the reader to [31] for a more in-depth discussion and exploratory data analysis of these 3 datasets. The first dataset uses MovieLens [27]. MovieLens is a dataset of individual ratings of movies, [31] turn it into a market by using matrix completion to fill in missing user-movie ratings, they then take the top 1500 most active users and 1500 most rated movies and set the valuations v_{ij} as the predicted ratings from the matrix completion. We also use the Jester Jokes dataset [26]. Here, we have 7200 individuals that have rated 100 jokes. We treat the jokes as the item to be allocated. Finally, we use the Household Items dataset introduced in [31]. Here we have 2876 survey takes entering a willingness to pay for 50 household items (vacuum cleaners, toasters, gas grills, etc.). For each dataset, we first rescale (w.l.o.g.) buyer valuations as described in §5.

We also consider an experiment on a simple infinite-dimensional market instance (which we refer to as "Inf-Dim") of $n = 100$ buyers and item space $\Theta = [0, 1]$, similar to the examples in [24, §4.2]. Let each buyer valuation v_i be normalized linear functions on $[0, 1]$, that is, $v_i(\theta) = c_i(\theta) + d_i$ such that $v_i(\Theta) := \int_\Theta v_i d\mu = \int_0^1 v_i(\theta) d\theta = 1 \Leftrightarrow \frac{c_i}{2} + d_i = 1$. We randomly generate (c_i, d_i) , $i = 1, \dots, n$ and run the dynamics for $T = 100n$ time steps.

For the finite dimensional datasets we compute equilibrium utilities u^* and utility prices β^* by solving the corresponding static instances using standard methods. For the infinite dimensional synthetic instance, we use the approach based on convex conic reformulation [24, §4] to compute β^* .

Figure 1 in §6 contains the plots for the MovieLens, Household Items and Inf-Dim datasets. Figure 2 contains the plots for the Jokes dataset.

Since items arrive one at a time, $t = 100$ time steps in a market with $n = 10$ buyers is very different from the same number of time steps in a market with $n = 1000$ buyers. To deal with this, we run PACE for $T = 100n$ time steps, referring to each n time steps as an *epoch*.

We record the average and maximum values of relative errors of the pacing multipliers β^t , time-averaged cumulative utilities \bar{u}^t and time-averaged expenditures \bar{b}^t .




Figure 2: Results for the same experiments as in Figure 1 (convergence of pacing multipliers, utilities and expenditures) on the Jokes dataset.

Convergence of expenditures to total budget. For each i , the quantity

$$\left| \frac{\bar{b}_i^t - B_i}{B_i} \right| = \left| \frac{\sum_{\tau=1}^t b_i^\tau - tB_i}{tB_i} \right|$$

can be viewed as the relative deviation of current cumulative expenditure at time t from the total budget tB_i available up to t . Hence, the residuals $\|(\bar{b}_i^t - B)/B\|/n$ and $\|(\bar{b}_i^t - B)/B\|_\infty$ are the average and maximum such deviations across all buyers. For each dataset (MovieLens, Household, Jokes and Inf-Dim), we plot the various quartiles of these residuals across all seeds, as shown in Figure 3.

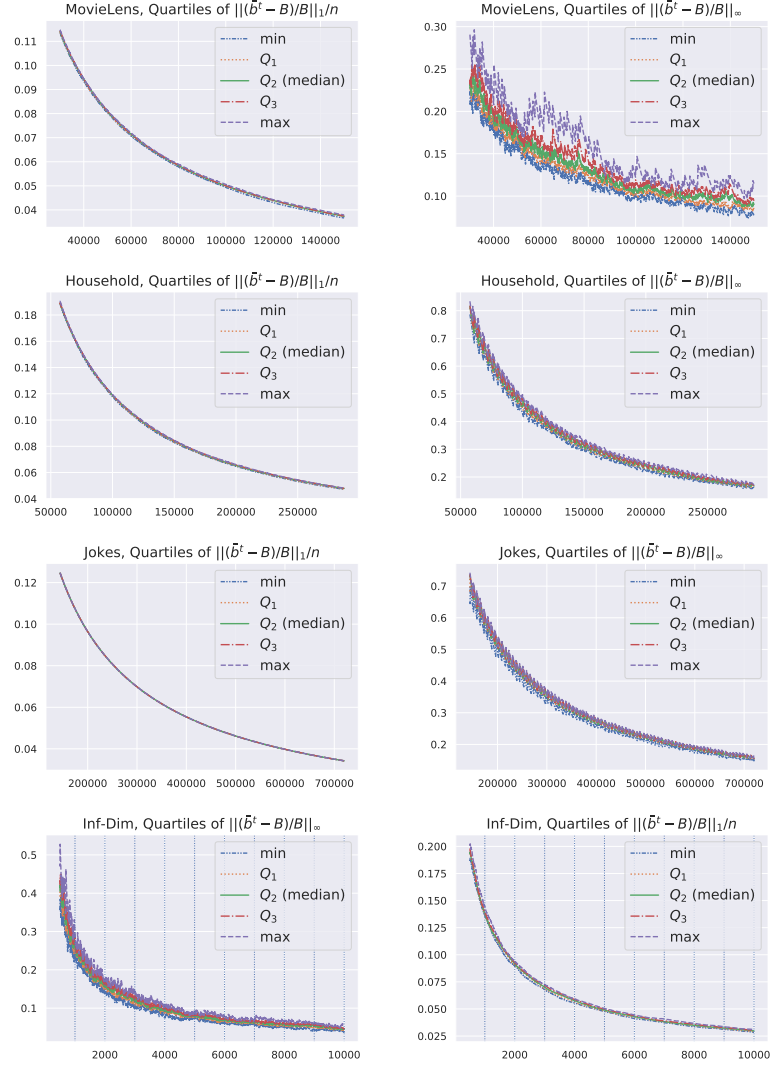


Figure 3: The PACE cumulative expenditure $\sum_{\tau=1}^t b_i^\tau$ of each buyer are close to the total amount of budget tB_i , as the quartile plots show. Vertical lines indicate when t is a multiple of $10n$.