

The Parity Ray Regularizer for Pacing in Auction Markets

Andrea Celli*
Bocconi University
Milan, Italy
andrea.celli2@unibocconi.it

Christian Kroer
Columbia University
New York, USA
christian.kroer@columbia.edu

Riccardo Colini-Baldeschi
Core Data Science, Meta
London, UK
rickuz@fb.com

Eric Sodomka
Core Data Science, Meta
London, UK
sodomka@fb.com

ABSTRACT

Budget-management systems are one of the key components of modern auction markets. Internet advertising platforms typically offer advertisers the possibility to pace the rate at which their budget is depleted, through *budget-pacing mechanisms*. We focus on *multiplicative pacing* mechanisms in an online setting in which a bidder is repeatedly confronted with a series of advertising opportunities. After collecting bids, each item is then allocated through a single-item, second-price auction. If there were no budgetary constraints, bidding truthfully would be an optimal choice for the advertiser. However, since their budget is limited, the advertiser may want to shade their bid downwards in order to preserve their budget for future opportunities, and to spread expenditures evenly over time. The literature on online pacing problems mostly focuses on the setting in which the bidder optimizes an additive separable objective, such as the total click-through rate or the revenue of the allocation. In many settings, however, bidders may also care about other objectives which oftentimes are non-separable. We study the frequent case in which the utility of a (proxy) bidder depends on the rewards obtained from items they are allocated, and on the distance of the realized distribution of impressions from a target distribution. We introduce a novel regularizer which can describe those distributional preferences, while keeping the problem tractable. We show that this regularizer can be integrated into an existing online mirror descent scheme with minor modifications, attaining the optimal order of sub-linear regret compared to the optimal allocation in hindsight when inputs are drawn independently, from an unknown distribution. Moreover, we show that our approach can easily be incorporated in *standard* existing pacing systems that are not usually built for this objective. The effectiveness of our algorithm in internet advertising applications is confirmed by numerical experiments on real-world data.

*Work done while at Meta (London, UK).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WWW '22, April 25–29, 2022, Virtual Event, Lyon, France

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9096-5/22/04...\$15.00

<https://doi.org/10.1145/3485447.3512061>

CCS CONCEPTS

• **Theory of computation** → **Computational pricing and auctions**.

KEYWORDS

budget-management systems, pacing mechanisms, online allocation

ACM Reference Format:

Andrea Celli, Riccardo Colini-Baldeschi, Christian Kroer, and Eric Sodomka. 2022. The Parity Ray Regularizer for Pacing in Auction Markets. In *Proceedings of the ACM Web Conference 2022 (WWW '22)*, April 25–29, 2022, Virtual Event, Lyon, France. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3485447.3512061>

1 INTRODUCTION

In the last decade, the spending on Internet advertising has grown dramatically, reaching more than \$130 billion in the United States in 2019 [23]. One of the drivers of this rapid growth has been the success of auction markets as a practical mechanism to match bidders (i.e., advertisers) to their target audience at an appropriate price [22, 40]. In this type of mechanisms, advertisers usually specify a targeting rule to target a certain population of users, and they are asked to bid on events of interest such as an impression, a click, a conversion or a video view. For each ad slot, the mechanism determines a winning advertiser, which is subsequently given the chance to show an impression and potentially generate the event of interest.

Auction markets are typically operationalized through simple mechanisms that can be effectively implemented at scale, and that simplify the interaction between advertisers and the platform. A common paradigm to manage large amounts of repeated ad auctions is for the platform to operate a *proxy bidder* on behalf of each advertiser. The advertisers communicate the maximum bid for the event of interest, their targeting criteria, and their overall budget for their campaigns to the proxy bidder. In real time, the proxy bidder then constructs bids for each individual auction via these parameters, as well as additional information computed by the platform.

The complexity of the competitive interactions taking place between advertisers, the frequency of the decisions, and the intrinsic uncertainties of the environment make budget management particularly challenging in this setting. An advertiser may lose out on significant amounts of revenue if their budget is depleted too early (thereby missing potentially valuable future bidding opportunities), or if their budget is not fully depleted within the planned

duration of the campaign. To solve this problem, platforms offer advertisers the possibility to *pace* the rate at which their budget is depleted, through a *budget-pacing mechanism*. The most common pacing mechanisms modify the values of the bids within a series of repeated auctions either by shading the bid appropriately (i.e., *multiplicative pacing*) [9, 18], or by determining a participation probability through the toss of an appropriately weighted coin for each auction (i.e., *probabilistic pacing*) [1, 9]. In this work, we focus on the former type of budget-management mechanism. This choice is motivated by the widespread adoption of this mechanisms in large online advertising platforms. Specifically, we study an online setting in which an advertiser is repeatedly confronted with a series of advertising opportunities, one per iteration, allocated through a second-price auction. If there were no budgetary constraints, bidding truthfully (i.e., bidding their true valuation for an advertising opportunity) would be an optimal choice for the advertiser. However, since their budget is limited, the advertiser may want to shade their bid downwards in order to preserve their budget for future opportunities.

The literature on online bidding and online pacing problems mostly focuses on settings in which the bidder optimizes an *additively separable objective*, such as the total click-through rate or revenue of the allocation (see, e.g., [10, 14, 15, 18, 26, 31]). In many settings, however, advertisers may also care about other objectives which oftentimes are *non-separable*, and therefore not amenable to traditional online learning techniques. Two notable exceptions are the work by Agrawal and Devanur [2], which focuses on solving general online stochastic convex programs with general concave objectives and convex constraints, and the recent work by Balseiro et al. [12], where the authors study online allocation problems with non-separable objectives in settings such as fairness across advertisers and load balancing.¹ We argue that another frequent objective, which is oftentimes sought after by advertisers, is reaching a certain distribution of impressions over a target population of users. Surprisingly enough, to the best of our knowledge, industry practitioners don't have at their disposal a practical way to solve this problem. Consider, for example, the following real-world use cases: an advertiser may want to enforce uniform sampling on a target population for conducting unbiased surveys, to promote a business based on a two-sided market by balancing their reach on users belonging to "both sides of the market" (this is the case, for example, of dating apps, and online marketplaces), or to perform online outreach to a population of users. In the latter setting, Gelauff et al. [28] showed that advertisers have to resort to complex segmentation strategies through subcampaigns to achieve a distribution of conversions close to that of the overall adult population of a city.

Our contributions. The above application scenarios are captured by the following utility model: the advertiser's utility is the sum of rewards accumulated during the repeated auctions, plus a penalty term which depends on the distance of the realized distribution of impressions from a target distribution, according to some given distance function. We consider distributional preferences over the advertising opportunities allocated to the advertiser, which are characterized by a feature vector determining their *category*. We consider a finite horizon model in which an advertiser participates

¹We devote Appendix A to a more in depth discussion of relevant previous works.

in a series of auctions. The advertiser has a fixed amount of non-replenishable resources (i.e., a budget). We consider an incomplete information model where, at each iteration, an item is drawn i.i.d. from some unknown distribution and presented to the advertiser, who has to compute a bid for the current item. Then, a second-price auction is run and the advertiser observes their reward and budget consumption. The advertiser does not get to observe the reward functions and types of future requests until their arrival. We propose a novel regularizer, which we call the *parity ray regularizer*, that can be employed to model the advertiser's distributional preferences according to some specific distance function. We show that, if the desired distance function satisfies some natural properties, the regularized online pacing problem can be addressed by casting the online mirror descent scheme by Balseiro et al. [12] to our setting. The algorithm guarantees regret of order $O(T^{1/2})$, which optimal in our stochastic setting. Moreover, we show that our framework could easily be incorporated into standard existing pacing systems that are not usually built for handling distributional preferences. Finally, we empirically demonstrate the efficacy of our mechanism on real-world data from a large Internet advertising company.

2 MULTIPLICATIVE PACING FOR AUCTION MARKETS

We study the problem faced by a bidder (i.e., an advertiser) who wishes to maximize their utility over a sequence of auctions that occur over T time steps. At each time $t = 1, \dots, T$ an item (i.e., an advertising opportunity to reach a user) appears, and an auction is run in order to allocate it. We will assume that the item is allocated via a single-item *second-price auction*, meaning that the highest bid wins, and the winner pays the second-highest bid. At every time t , the bidder observes their valuation v_t for the item, submits a bid b_t , and then observes the variable $x_t \in \{0, 1\}$ specifying whether they won, and if they did win then they observe the price p_t that they paid for the item.

Without further restrictions, bidding in such a dynamic auction setting would be straightforward: the bidder can simply bid their true value, and the second-price auction ensures that this maximizes their utility. However, in practice, bidders usually have a total budget B , and they wish to constrain their expenditure across the T auctions to be at most B . This complicates bidding even in second-price auctions, because this budget constraint links all T auctions together.² If we allow fractional allocations, then an omniscient budget-constrained bidder can be thought of as a fractional knapsack packer: they now wish to order the T auctions by *decreasing bang-per-buck* v_t/p_t , and grab items in this order until they either deplete their budget, or they reach some item t such that $v_t/p_t \leq 1$ [18]. In practice, the bidder must decide their bids in an online fashion: at time t they only know their first t valuations (v_1, \dots, v_t) , their past allocations up to time $t - 1$, i.e., (x_1, \dots, x_{t-1}) , as well as their past expenditures $(x_1 p_1, \dots, x_{t-1} p_{t-1})$. Thus, the only information they have about the auction at time t is their valuation for the item, and they must adaptively smooth out their

²Here, we assume that the bidder cannot manipulate future prices by bidding higher when they lose, in order to encourage the depletion of the budget of other bidders. This assumption is reasonable in large-scale markets, as we will discuss in Section 3.

spending across the T auctions, in order to satisfy their budget constraint while maximizing utility.

The above problem of how to bid in such a setting has been studied extensively both from the perspective of an individual bidder, as well as from a platform perspective. One popular approach in practice is to use a *pacing multiplier* $\alpha \in [0, 1]$, which is then used to construct each bid as $b_t = \alpha v_t$. Given a sequence of prices (p_1, \dots, p_T) , an optimal allocation can be achieved within the context of second-price auctions by choosing α such that the bidder spends their budget exactly (assuming that the bidder can choose the fraction that they wish to win on any items such that $\alpha v_t = p_t$) [13, 18]. The pacing multiplier can be adaptively controlled via online-learning methods in order to guarantee asymptotic optimality in stochastic environments [10, 14].

In the above, we presented the setting as if the individual bidder controls their pacing parameter α used in smoothing their expenditure. However, a second model which arises in practice is the *proxy-bidder setting*. The proxy-bidder setting arises in online advertising, where advertisers often submit only a value per click, budget, and targeting criteria to the platform. The platform then performs budget smoothing on behalf of the advertiser. In that setting, the proxy bidder refers to a bidder operated by the platform, which is effectively a control algorithm that scales the parameter α up or down, depending on whether the advertiser is on track to spend their budget correctly across the given time frame. Equilibria arising in the proxy-bidder setting have been studied in several works (see, e.g., [9, 17, 18]).

3 PROBLEM DESCRIPTION

In this paper, we focus on the setting where we wish to smooth out the expenditure of a single bidder, while the rest of the market behaves according to a stochastic model. In particular, we will assume that valuations, prices, and item categories are generated i.i.d. from an unknown distribution. In practice, valuations and categories can be safely described as coming from a stochastic model. The prices, however, typically come from the paced bids arising from other bidders whom are also employing some budget-smoothing technique. In small-scale settings where bidders react to the expenditure of other bidders, these prices could behave adversarially rather than stochastically. However, for large-scale markets, an individual bidder has almost no impact on the prices of the overall market, in which case stochastic behavior of the prices is a reasonable assumption. This idea can be formalized, for example, via fluid mean-field models such as the one given by Balseiro et al. [13].

The bidder can spend a budget B over a sequence of auctions within a finite time horizon T . Let $x \in [0, 1]^T$ be the allocation vector for the bidder, and $x_t \in \mathcal{X}_t \subseteq [0, 1]$ be the allocation at time t . Let $\rho \in \mathbb{R}_{>0}$ be the *per-iteration budget*, which is a constant such that $B = T\rho$. Items can be of m different *types* or, equivalently, *categories*. The vector $c_t \in [0, 1]^m$ specifies the categories of the item being auctioned at time t . At each iteration t , the bidder receives as input the tuple (v_t, p_t, c_t) specifying the valuation for the item $v_t \in \mathbb{R}_{\geq 0}$, its price $p_t \in \mathbb{R}_{\geq 0}$, and its category vector c_t . At each t , the input tuple is generated i.i.d. from an unknown distribution $\mathcal{P} \in \Delta^{\mathcal{I}}$, where \mathcal{I} is the set of all possible input tuples that can be

observed.³ In the remainder of the paper we make the following assumption on the regularity of inputs, which is easily satisfied in practice.

ASSUMPTION 1. *There exists $\bar{v} \in \mathbb{R}_{\geq 0}$ and $\bar{p} \in \mathbb{R}_{\geq 0}$ such that, for any input tuple $(v_t, p_t, c_t) \in \mathcal{I}$ in the support of \mathcal{P} , it holds $v_t \leq \bar{v}$ and $p_t \leq \bar{p}$.*

Let $f_t : [0, 1] \rightarrow \mathbb{R}$ be the reward function at time t such that, for each $x \in \mathcal{X}_t$,

$$f_t(x) := (v_t - p_t)x. \quad (1)$$

In the standard model of bidder's utility (i.e., when distributional preferences on categories are disregarded), the bidder is interested in finding an online algorithm guaranteeing performance *close* to that of the optimal solution in hindsight. In particular, this baseline is the expected reward that the bidder would achieve by computing the optimal allocation x when the input sequence from $t = 1$ to T is known in advance. This amounts to solving the offline optimal allocation problem for each possible realized input sequence, and by taking expectations according to the realization probability specified by \mathcal{P} . Formally, the value of the optimal allocation is defined as follows

$$\mathbb{E}_{\mathcal{P}} \left[\max_{x: x_t \in \mathcal{X}_t} \sum_{t=1}^T f_t(x_t) \text{ s.t. } \sum_{t=1}^T p_t x_t \leq T\rho \right]. \quad (2)$$

4 PARITY RAY REGULARIZER

The question now is whether we can endow Equation (2) with a practical way of taking bidder's distributional preferences into account. For simplicity, we assume that a type specifies a convex combination over categories, that is, for any t , the vector of realized categories c_t belongs to $C := \{c \in [0, 1]^m : \sum_i c_i = 1\}$. This assumption can be relaxed to a larger bounded set. Given a target distribution over categories $\hat{x} \in \Delta^m$, we consider a setting where the bidder is interested in steering their realized distribution of impressions towards \hat{x} according to some distance function D . Given a time horizon T , a sequence of types $(c_t)_{t=1}^T$, and $\hat{x} \in \Delta^m$, the bidder's goal is to obtain a vector of allocations $x \in [0, 1]^T$ such that, for each $i \in [m]$, $\sum_{t=1}^T c_{t,i} x_t / \sum_{t=1}^T x_t$ is *close* to \hat{x}_i according to D . However, directly handling such realized distribution of impressions would yield a number of significant technical difficulties, as shown in the following example.

Example 4.1. Let us focus on the simple case in which the bidder chooses the ℓ_∞ norm as the desired distance function D . Then, the bidder could determine a fixed value $\alpha \in \mathbb{R}_{\geq 0}$, according to their distributional preferences, and add the following additional constraint to Equation (2)

$$\left\| \frac{\sum_{t=1}^T c_t x_t}{\sum_{t=1}^T x_t} - \hat{x} \right\|_\infty \leq \alpha,$$

³The set $\{1, \dots, n\}$, where $n \in \mathbb{N}_{>0}$, is compactly denoted as $[n]$; the empty set as \emptyset . Given a set \mathcal{X} , we denote its convex hull with the symbol $\text{co } \mathcal{X}$. Vectors are marked in bold. We denote by $\langle \cdot, \cdot \rangle$ the scalar product between two vectors. Given a discrete set \mathcal{X} , we denote by $\Delta^{\mathcal{X}}$ the $|\mathcal{X}|$ -simplex, that is, the set $\Delta^{\mathcal{X}} := \{v \in \mathbb{R}_{\geq 0}^{|\mathcal{X}|} : \sum_i v_i = 1\}$. Analogously, we denote the *full-dimensional standard simplex* over \mathcal{X} by $\Delta_+^{\mathcal{X}} := \{v \in \mathbb{R}_{\geq 0}^{|\mathcal{X}|} : \sum_i v_i \leq 1\}$. The symbols Δ^n and Δ_+^n , with $n \in \mathbb{N}_{>0}$, are used to denote $\Delta^{[n]}$ and $\Delta_+^{[n]}$, respectively.

that is, $|\sum_{t=1}^T c_{t,i}x_t / \sum_{t=1}^T x_t - \hat{x}_i| \leq \alpha$ for each $i \in [m]$. Each of the latter constraints could be suitably linearized. However, this would yield a set of constraints which are not *packing*, that is they are not of the form $\sum_t b_t x_t \leq K$ for some $b_t \geq 0$, and $K \geq 0$. Indeed, after linearizing and taking variables to the left-hand side, there would be several negative coefficients in front of variables. However, the dual mirror descent framework—which is the standard algorithm employed for online allocation in this setting—is only set up to work with packing constraints (see, e.g., [10, 12]).

Another natural option would be embedding an additive penalty term directly into the objective function of Equation (2). In particular, one could consider the following objective function

$$\max_{\mathbf{x}: \mathbf{x}_t \in \mathcal{X}_t} \sum_{t=1}^T f_t(x_t) - \alpha D\left(\hat{\mathbf{x}}, \frac{\sum_{t=1}^T \mathbf{c}_t x_t}{\|\mathbf{x}\|_1}\right),$$

for some $\alpha \in \mathbb{R}_{\geq 0}$. However, we observe that if the distance function takes directly in input the current distribution of impressions and the target one, then solving the new version of Equation (2) would entail solving a non-convex optimization problem due to normalization.

Then, a natural question is the following: *is it possible to incorporate such distributional preferences as a concave regularization term in the bidder's objective function?* We answer in the positive by proposing the following concave regularization term as a natural solution for this problem.

Definition 4.2 (Parity ray regularizer). Let $D : \Delta_+^m \times \Delta_+^m \rightarrow \mathbb{R}_{\geq 0}$. Then, given a target distribution $\hat{\mathbf{x}} \in \Delta^m$, the *parity ray regularizer* is a function $R : \Delta_+^m \times \Delta_+^m \rightarrow \mathbb{R}_{\leq 0}$ such that, for any vector $\bar{\mathbf{x}} \in \Delta_+^m$, it holds

$$R(\bar{\mathbf{x}}) := - \min_{\gamma \in [0,1]} D(\gamma \hat{\mathbf{x}}; \bar{\mathbf{x}}).$$

Intuitively, D acts as a pseudo-distance measure on Δ_+^m . It is not a true distance metric because we will not be assuming that it satisfies all the conditions of a metric. It can be thought of as analogous to a divergence from statistics, although we measure distances on the full-dimensional simplex. Given some current distribution $\bar{\mathbf{x}}$, R performs a projection onto the line segment generated by $[0, 1] \times \hat{\mathbf{x}}$, where the projection is in terms of D (this is analogous to, e.g., a Bregman projection).

The “parity ray” name comes from the fact that we can think of R as measuring the distance between a given point $\bar{\mathbf{x}} \in \Delta_+^m$ to the ray $\{\gamma \hat{\mathbf{x}} : \gamma \geq 0\}$. In our definition we only consider a segment of this ray by restricting $\gamma \in [0, 1]$. This restriction is mostly for analytical convenience when deriving Lipschitz constants below. For a suitable D that extends to the positive orthant, one could remove the restriction and measure distances on the whole positive orthant.

We will assume that D satisfies a number of properties that will allow us to optimize it as part of an online optimization procedure.

ASSUMPTION 2. *We assume that $D(\mathbf{x}, \bar{\mathbf{x}})$ satisfies the following conditions:*

- D is jointly convex on $\Delta_+^m \times \Delta_+^m$, and $D(\mathbf{x}; \bar{\mathbf{x}}) \in [0, \bar{r}]$ for all $\mathbf{x}, \bar{\mathbf{x}} \in \Delta_+^m$.
- For a fixed γ , the function $D(\gamma \hat{\mathbf{x}}; \cdot)$ is L -Lipschitz with respect to the second argument under some ℓ_p norm.

PROPOSITION 4.3. *Given Assumption 2, the regularizer R is concave, closed, and for each $\mathbf{x} \in \Delta_+^m$, $R(\mathbf{x}) \in [\underline{r}, 0]$ with $\underline{r} = -\bar{r}$. Furthermore, R is L -Lipschitz continuous under the ℓ_p norm.*

PROOF. Concavity follows from the facts that D is jointly convex and convexity is preserved by taking the infimum over one argument, finally taking the negative leads to concavity. The fact that R is L Lipschitz follows since D is L Lipschitz in the second argument for each γ , and taking the infimum preserves Lipschitz continuity when all functions in the infimum have the same Lipschitz constant. Closedness is implied by L Lipschitzness, since closedness of a convex function is equal to lower-semicontinuity, a weaker condition than Lipschitz continuity. \square

There are various choices of D that comply with Assumption 2, we mention the following notable examples.

- For any ℓ_p norm, the distance $D_p(\mathbf{x}; \bar{\mathbf{x}}) = \|\mathbf{x} - \bar{\mathbf{x}}\|_p$ satisfies Assumption 2. In particular, D_p is 1 Lipschitz with respect to ℓ_p , and the remaining properties follow directly from the definition of a norm and the fact that D_p is defined on a bounded set Δ_+^m . Moreover, the ℓ_2 norm yields a simple closed form solution for the parity ray regularizer, which is particularly appealing for practical purposes. In particular, we have that for any $\bar{\mathbf{x}} \in \Delta_+^m$, $R(\bar{\mathbf{x}}) = -\|\bar{\mathbf{x}} - (\langle \bar{\mathbf{x}}, \hat{\mathbf{x}} \rangle / \|\hat{\mathbf{x}}\|_2^2) \hat{\mathbf{x}}\|_2$, which is well defined with respect to Definition 4.2 since $0 \leq \langle \bar{\mathbf{x}}, \hat{\mathbf{x}} \rangle / \|\hat{\mathbf{x}}\|_2^2 \leq 1$.
- Another source of suitable choices for D are f -divergences: $D_f(\mathbf{x}; \bar{\mathbf{x}}) = \sum_{i=1}^m \bar{x}_i f(x_i / \bar{x}_i)$ defined with respect to some convex function f , since f -divergences are jointly convex. However, f -divergences are not necessarily L Lipschitz for a fixed L . In the case where we wish to use a non-Lipschitz f , we have to modify f such that $|\frac{s}{t^2} f'(s/t)| \leq L$ for $s, t \in [0, 1]$. For example, setting $f(t) = t \log t$ yields the Kullback-Leibler (KL) divergence. However, we need to slightly modify the KL divergence in order to ensure that Assumption 2 is satisfied: (1) we must shift it by $\log m$ to make it nonnegative, and (2) we make it Lipschitz in the second argument by setting $f(t) = t \log \min(t, 1/\epsilon)$; ϵ then can be interpreted as a lower probability threshold at which we stop increasing the penalty. This lower threshold may actually be useful in practice, to prevent the algorithm from reacting too aggressively to allocation imbalances at the first few iterations.

Finally, equipped with Definition 4.2, computing the offline optimal allocation now amounts to the following convex optimization problem

$$\text{OPT}(\mathcal{P}) := \mathbb{E}_{\mathcal{P}} \left[\max_{\mathbf{x}: \mathbf{x}_t \in \mathcal{X}_t} \sum_{t=1}^T f_t(x_t) + TR \left(\frac{1}{T} \sum_{t=1}^T \mathbf{c}_t x_t \right) \right] \quad \text{s.t.} \quad \sum_{t=1}^T p_t x_t \leq T\rho \quad (3)$$

We say that a bidder is optimizing for their *parity-regularized utility* if they are solving Problem (3) and R is defined as in Definition 4.2.

5 ONLINE ALLOCATION ALGORITHM

First, we show that a suitably defined Lagrangian dual function provides an upper bound on $\text{OPT}(\mathcal{P})$ (omitted proofs are reported

in Appendix B). Let $f_t^* : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ be the conjugate function of f_t as defined in Equation (1) restricted to \mathcal{X}_t , that is, for all $\mu \in \mathbb{R}_{\geq 0}$,

$$f_t^*(\mu) := \sup_{x \in \mathcal{X}_t} \{f_t(x) - \mu x\}.$$

Moreover, let $R^* : \mathbb{R}^m \rightarrow \mathbb{R}$ be the conjugate function of the regularizer R restricted to Δ_+^m . Formally, for each $\lambda \in \mathbb{R}^m$,

$$R^*(-\lambda) := \sup_{\bar{x} \in \Delta_+^m} \{R(\bar{x}) + \langle \lambda, \bar{x} \rangle\}.$$

Let \mathcal{D} be the set of dual variables for which the conjugate of the regularizer is bounded. Formally,

LEMMA 5.1. *It holds*

$$\mathcal{D} := \{(\mu, \lambda) \in \mathbb{R}_{\geq 0} \times \mathbb{R}^m : -\infty < R^*(-\lambda) < +\infty\} = \mathbb{R}_{\geq 0} \times \mathbb{R}^m.$$

For a given distribution \mathcal{P} over the finite set of possible input tuples \mathcal{I} , the Lagrangian dual function $D(\mu, \lambda | \mathcal{P}) : \mathcal{D} \rightarrow \mathbb{R}$ is such that, for each pair of dual variables $(\mu, \lambda) \in \mathcal{D}$,

$$D(\mu, \lambda | \mathcal{P}) := \mathbb{E}_{(v_t, p_t, c_t) \sim \mathcal{P}} [f_t^*(\mu p_t + \langle \lambda, c_t \rangle) + R^*(-\lambda)] + \rho \mu. \quad (4)$$

Then, the following holds.

THEOREM 5.2. *Given a time horizon $T \geq 0$, for any $(\mu, \lambda) \in \mathcal{D}$, $OPT(\mathcal{P}) \leq TD(\mu, \lambda | \mathcal{P})$.*

Given this intermediate result, we describe an online allocation algorithm (Algorithm 1) based on the online mirror descent scheme by Balseiro et al. [10], which we adapt to account for parity-regularized objectives (which are non-separable). We exploit the specific structure of our problem to recover better constants in the regret upper bound (see Theorem 5.4). For a specific choice of the reference function ψ , Algorithm 1 gives an instantiation of the dual subgradient descent algorithm by Balseiro et al. [12] for our setting.

The algorithm proceeds according to the following main steps:

Primal decision At each time step t , given input tuple (v_t, p_t, c_t) , the algorithm computes the optimal allocation x^* which maximizes an opportunity cost-adjusted reward, based on the current dual solutions (μ_t, λ_t) (Equation 5). The optimal decision x^* is taken if the remaining budget is enough to cover costs, that is, $p_t x^* \leq B_t$. Finally, the target distribution over categories \bar{x}_t is computed by computing the value maximizing the regularizer R adjusted by an additive term accounting for the current dual solution λ_t (Equation 6).

Dual variables update First, the algorithm computes an unbiased stochastic estimator of a subgradient of $D(\mu, \lambda | \mathcal{P})$ at (μ_t, λ_t) (Equation 7). The algorithm employs this estimator to update the vector of dual variables by performing an online dual mirror descent step with step size η and reference function ψ (Equation 8).

Intuitively, at each time step t , the algorithm compares (i) the actual expenditure from the auction to the expected rate of spend per iteration ρ , and (ii) the target distribution over categories at t to the realized category. If the actual expenditure is higher (resp., lower) than this expected rate, then the algorithm surmises that future opportunities will offer higher (resp., lower) *bang for the buck* than current opportunities, and therefore the algorithm lowers (resp., raises) the bid shading multiplier. At the same time, if the realized category causes an undesired skew, and moves the realized distribution of impressions away from the desired \bar{x}_t , then the

Algorithm 1 Online algorithm for parity-regularized pacing

1: **Input:** initial dual solutions (μ_1, λ_1) , time horizon T , initial budget consumption $B_1 = B$, function ψ , step-size η

2: **for** $t = 1, \dots, T$ **do**

3: Observe $(v_t, p_t, c_t) \sim \mathcal{P}$

4: Compute primal decision:

5:

$$x^* \in \arg \max_{x \in \mathcal{X}_t} \{f_t(x) - \mu_t p_t x - \langle \lambda_t, c_t x \rangle\} \quad (5)$$

6:

$$x_t = \begin{cases} x^* & \text{if } p_t x^* \leq B_t \\ 0 & \text{otherwise} \end{cases}$$

7:

$$\bar{x}_t \in \arg \max_{\bar{x}' \in \Delta_+^m} \{R(\bar{x}') + \langle \lambda_t, \bar{x}' \rangle\} \quad (6)$$

8:

9: Update resource consumption: $B_{t+1} = B_t - p_t x_t$

10: Compute subgradient $g \in \mathbb{R}^{m+1}$ of $D(\mu_t, \lambda_t | \mathcal{P})$:

11:

$$g = (\rho - p_t x^*, \bar{x}_t - c_t x^*) \quad (7)$$

12:

13: Update dual variables:

14:

$$(\mu_{t+1}, \lambda_{t+1}) = \arg \min_{(\mu, \lambda) \in \mathcal{D}} \left\{ \langle g, (\mu, \lambda) \rangle + \frac{1}{\eta} B_{\psi}((\mu, \lambda), (\mu_t, \lambda_t)) \right\} \quad (8)$$

algorithm will adjust its dual solutions to penalize allocations from that category, and increase the likelihood of acceptance for items coming from *under-represented* categories.

5.1 Regret bound

We focus on the stochastic setting in which, at each t , an input tuple is drawn i.i.d. from \mathcal{P} . This setting is particularly relevant for applications to large Internet advertising platforms in which the number of bidders interacting makes the environment unlikely to react adversarially to choices of one specific bidder. We show that Algorithm 1 attains a sublinear regret of order $O(T^{1/2})$.

Let τ be the time at which the budget of the bidder is depleted when following Algorithm 1. By playing according to Algorithm 1, the bidder attains the following expected reward:

$$\text{REW}(\mathcal{P}) := \mathbb{E}_{\mathcal{P}} \left[\sum_{t=1}^T f_t(x_t) + TR \left(\frac{1}{T} \sum_{t=1}^T c_t x_t \right) \right]. \quad (9)$$

We start with the following intermediate result which provides a lower bound on the expected reward attained by following Algorithm 1 up to the time τ in which the budget is fully depleted.

LEMMA 5.3. *Consider an arbitrary $\mathcal{P} \in \Delta^{\mathcal{I}}$. For each $t \in [T]$, let x_t and \bar{x}_t be computed according to Equation (5) and Equation (6), respectively. Moreover, given the time of budget depletion τ , let $\bar{\mu} =$*

$1/\tau \sum_{t=1}^{\tau} \mu_t$ and $\bar{\lambda} = 1/\tau \sum_{t=1}^{\tau} \lambda_t$. Then, it holds

$$\mathbb{E}_{\mathcal{P}} \left[\sum_{t=1}^{\tau} (f_t(x_t) + R(\bar{x}_t)) \right] \geq \mathbb{E}_{\mathcal{P}} \left[\tau D(\bar{\mu}, \bar{\lambda} | \mathcal{P}) - \sum_{t=1}^{\tau} (\mu_t(\rho - p_t x_t) - \langle \lambda_t, \bar{x}_t - c_t x_t \rangle) \right].$$

In order to prove the regret bound of Theorem 5.4 we need the following additional assumption on the reference function ψ of the Bregman divergence B_{ψ} .

ASSUMPTION 3. *The reference function ψ is σ -strongly convex with respect to an ℓ_p norm $\|\cdot\|_p$.*

This assumption is a relaxation of the requirements in Balseiro et al. [10, Assumption 2] (i.e., strong convexity with respect to $\|\cdot\|_1$). Indeed, by exploiting the specific structure of our problem, we can prove the following.

THEOREM 5.4. *Consider Algorithm 1 with step-size $\eta \geq 0$ and initial solution $(\mu_1, \lambda_1) \in \mathcal{D}$. Suppose Assumption 1, Proposition 4.3, and Assumption 3 are satisfied, and the requests are drawn i.i.d. from an unknown distribution $\mathcal{P} \in \Delta^I$. Then, for any $T \geq 1$, it holds*

$$\text{OPT}(\mathcal{P}) - \text{REW}(\mathcal{P}) \leq C_1 + \frac{G^2 \eta}{\sigma} T + \frac{1}{\eta} C_2,$$

where $C_1 = (\bar{v} - \underline{r} + 2L)\bar{p}/\rho$, $G = \max\{\rho + \bar{p}, 2\}$, and

$$C_2 = 2 \sup \{B_{\psi}((\mu, \lambda), \mathbf{d}_1) : (\mu, \lambda) \in \mathcal{D}, \|\lambda\| \leq L\}.$$

By setting $\eta \sim T^{-1/2}$ we obtain a regret of order $O(T^{1/2})$, which is the optimal order of regret that can be attained in this setting as shown by Arlotto and Gurvich [6, Lemma 1].

PROOF. Recall that τ is the time at which the budget is depleted. By Theorem 5.2 and since $v_t \leq \bar{v}$ and $p_t \geq 0$ we have that, for any $(\mu, \lambda) \in \mathcal{D}$,

$$\begin{aligned} \text{OPT}(\mathcal{P}) &= \frac{\tau}{T} \text{OPT}(\mathcal{P}) + \frac{T-\tau}{T} \text{OPT}(\mathcal{P}) \\ &\leq \tau D(\mu, \lambda | \mathcal{P}) + (T-\tau)\bar{v}. \end{aligned}$$

Let $\xi_{1,t}(\mu) = \mu(\rho - p_t x_t)$ and $\xi_{2,t}(\lambda) = \langle \lambda, \bar{x}_t - c_t x_t \rangle$ be the two complementary slackness terms at iteration t . From Equation (9) and Lemma 5.3 we have

$$\begin{aligned} \text{REW}(\mathcal{P}) &= \mathbb{E} \left[\sum_{t=1}^{\tau} f_t(x_t) + TR \left(\frac{1}{T} \sum_{t=1}^T c_t x_t \right) \right] \\ &\geq \mathbb{E} \left[\tau D(\bar{\mu}, \bar{\lambda} | \mathcal{P}) - \sum_{t=1}^{\tau} (\xi_{1,t}(\mu_t) + \xi_{2,t}(\lambda_t) + R(\bar{x}_t)) \right. \\ &\quad \left. + TR \left(\frac{1}{T} \sum_{t=1}^T c_t x_t \right) \right]. \end{aligned}$$

Then,

$$\begin{aligned} \text{OPT}(\mathcal{P}) - \text{REW}(\mathcal{P}) &\leq \mathbb{E} \left[(T-\tau)\bar{v} \right. \\ &\quad \left. + \sum_{t=1}^{\tau} (\xi_{1,t}(\mu_t) + \xi_{2,t}(\lambda_t) + R(\bar{x}_t)) - TR \left(\frac{1}{T} \sum_{t=1}^T c_t x_t \right) \right]. \quad (10) \end{aligned}$$

Complementary slackness term $\xi_{1,t}$. For each t , the gradient of $\xi_{1,t}$ is given by $\nabla_{\mu} \xi_{1,t} = \rho - p_t x_t$ and is bounded as follows $|\nabla_{\mu} \xi_{1,t}| \leq \rho + \bar{p}$. Algorithm 1 applies online mirror descent to the sequence of functions $(\xi_{1,t}, \xi_{2,t})_{t=1}^T$. Therefore, by holding the second term fixed, we have that, for any pair $(\mu, \lambda) \in \mathcal{D}$, $\sum_{t=1}^{\tau} \xi_{1,t}(\mu_t) - \xi_{1,t}(\mu) \leq U(\tau, (\mu, \lambda)) \leq U(T, (\mu, \lambda))$, where $U(\tau, (\mu, \lambda))$ is the regret guarantee for online mirror descent after τ iterations, and the second inequality follows from the fact the regret upper bound is increasing in t .

Complementary slackness term $\xi_{2,t}$. For each t , the gradient of $\xi_{2,t}$ is given by $\nabla_{\lambda} \xi_{2,t} = \bar{x}_t - c_t x_t$. Since $\bar{x}_t, c_t x_t \in \Delta_+^m$, the gradient is bounded by $\|\nabla_{\lambda} \xi_{2,t}\|_{p,*} \leq 2$. Therefore, analogously to the previous case, for any $(\mu, \lambda) \in \mathcal{D}$, we have $\sum_{t=1}^{\tau} \xi_{2,t}(\lambda_t) - \xi_{2,t}(\lambda) \leq U(T, (\mu, \lambda))$. Now, we focus on a particular choice of λ . Let

$$\hat{\lambda} \in \arg \max_{\lambda \in \mathbb{R}^m} \left\{ R^*(-\lambda) - \langle \lambda, \frac{1}{T} \sum_{t=1}^T c_t x_t \rangle \right\}. \quad (11)$$

By definition of R^* we have that, for any t ,

$$R^*(-\lambda) := \max_{\bar{x} \in \Delta_+^m} \{R(\bar{x}) + \langle \lambda, \bar{x} \rangle\} \geq R(\bar{x}_t) + \langle \lambda, \bar{x}_t \rangle.$$

Then, by summing over $t = 1, \dots, T$ we obtain

$$R^*(-\lambda) \geq \frac{1}{T} \sum_{t=1}^T (R(\bar{x}_t) + \langle \lambda, \bar{x}_t \rangle). \quad (12)$$

Now, if R is concave and closed we have that $R = R^{**}$ (see, e.g., Borwein and Lewis [16, Theorem 4.2.1]). Then, by using this fact and Equation (11),

$$R \left(\frac{1}{T} \sum_{t=1}^T c_t x_t \right) = R^{**} \left(\frac{1}{T} \sum_{t=1}^T c_t x_t \right) = R^*(-\hat{\lambda}) - \langle \hat{\lambda}, \frac{1}{T} \sum_{t=1}^T c_t x_t \rangle.$$

Then,

$$R^*(-\hat{\lambda}) = R \left(\frac{1}{T} \sum_{t=1}^T c_t x_t \right) + \langle \hat{\lambda}, \frac{1}{T} \sum_{t=1}^T c_t x_t \rangle. \quad (13)$$

By Equation (12) and Equation (13) we obtain

$$R \left(\frac{1}{T} \sum_{t=1}^T c_t x_t \right) + \langle \hat{\lambda}, \frac{1}{T} \sum_{t=1}^T c_t x_t \rangle \geq \frac{1}{T} \sum_{t=1}^T (R(\bar{x}_t) + \langle \hat{\lambda}, \bar{x}_t \rangle).$$

The above inequality can be rewritten as

$$\sum_{t=1}^T \xi_{2,t}(\hat{\lambda}) = \sum_{t=1}^T \langle \hat{\lambda}, \bar{x}_t - c_t x_t \rangle \leq R \left(\sum_{t=1}^T c_t x_t \right) - \sum_{t=1}^T R(\bar{x}_t). \quad (14)$$

Finally, by setting $\hat{\lambda}$ as specified in Equation (11), and for any $\mu \in \mathbb{R}_{\geq 0}$, we can bound the complementary slackness term as follows

$$\begin{aligned} \sum_{t=1}^{\tau} \xi_{2,t}(\lambda_t) &\leq \sum_{t=1}^{\tau} \xi_{2,t}(\hat{\lambda}) + U(T, (\mu, \hat{\lambda})) \\ &= \sum_{t=1}^{\tau} \xi_{2,t}(\hat{\lambda}) - \sum_{t=\tau+1}^T \xi_{2,t}(\hat{\lambda}) + U(T, (\mu, \hat{\lambda})) \\ &\leq R \left(\sum_{t=1}^T c_t x_t \right) - \sum_{t=1}^T R(\bar{x}_t) - \sum_{t=\tau+1}^T \xi_{2,t}(\hat{\lambda}) + U(T, (\mu, \hat{\lambda})), \end{aligned}$$

where the last inequality follows from Equation (14).

Final step. By substituting in Equation (10) the above upper bounds to the complementary slackness terms we obtain that, for any $\mu \in \mathbb{R}_+$ and for $\hat{\lambda}$ as in Equation (11),

$$\begin{aligned} \text{OPT}(\mathcal{P}) - \text{REW}(\mathcal{P}) &\leq \mathbb{E} \left[(T - \tau) \bar{v} + \sum_{t=1}^{\tau} \xi_{1,t}(\mu) \right. \\ &\quad \left. - \sum_{t=\tau+1}^T \left(R(\bar{x}_t) + \xi_{2,t}(\hat{\lambda}) \right) \right] + 2U(T, (\mu, \hat{\lambda})) \\ &\leq \mathbb{E} \left[(T - \tau) (\bar{v} - \underline{r} + 2\|\hat{\lambda}\|_{\infty}) + \sum_{t=1}^{\tau} \xi_{1,t}(\mu) \right] + 2U(T, (\mu, \hat{\lambda})). \end{aligned}$$

By definition of the stopping time τ , $\sum_{t=1}^{\tau} p_t x_t + \bar{p} \geq T\rho$. By letting $C = (\bar{v} - \underline{r} + 2\|\hat{\lambda}\|_{\infty})$, and $\hat{\mu} = C/\rho$, which is well defined as $\hat{\mu} \geq 0$, we obtain

$$\sum_{t=1}^{\tau} \xi_{1,t}(\hat{\mu}) = \sum_{t=1}^{\tau} \hat{\mu}(\rho - p_t x_t) \leq \hat{\mu}(\tau\rho - T\rho + \bar{p}) = \frac{C}{\rho} \bar{p} - (T - \tau)C.$$

Then, $\text{OPT}(\mathcal{P}) - \text{REW}(\mathcal{P}) \leq C\bar{p}/\rho + 2U(T, (\mu, \lambda))$. Finally, we observe that $\|\hat{\lambda}\|_{\infty} \leq L$ because R is L -Lipschitz continuous with respect to $\|\cdot\|_1$ and $\hat{\lambda} \in \partial R(\bar{x}_t)$ by Equation (11). Therefore, by substituting the standard regret guarantees for OMD (see, e.g., Orabona [39, Section 6]), we have

$$\text{OPT}(\mathcal{P}) - \text{REW}(\mathcal{P}) \leq C_1 + \frac{G^2\eta}{\sigma} T + \frac{1}{\eta} C_3,$$

where $C_1 = (\bar{v} - \underline{r} + 2L)\bar{p}/\rho$, $G = \max\{\rho + \bar{p}, 2\}$, and

$$C_3 = 2 \sup\{B_{\psi}((\mu, \lambda), (\mu_1, \lambda_1)) : (\mu, \lambda) \in \mathcal{D}, \|\lambda\| \leq L\}.$$

This concludes the proof. \square

6 REGULARIZED PACING WITHIN A PACED-AUCTION FRAMEWORK

One of the major selling points of our approach to enforcing distributional preferences via regularization is that Algorithm 1 can easily be incorporated in existing pacing systems that are not necessarily built around this objective. In Section 2 we described how pacing systems with proxy bidders work for large-scale Internet advertising. In that setup, the platform controls a proxy bidder on behalf of every advertiser i , where the proxy bidder uses a control algorithm on the pacing parameter α_i used to construct bids, in order to ensure that the advertiser satisfies their budget constraint. An important feature of this setup is that each proxy bidder can control their own pacing parameter α_i purely through their observed spending. This leads to a highly decentralized framework, where communication between proxy bidders is only necessary via the second-price auctions that sell items whenever they show up, and bidders only need to submit their paced bids to these auctions. This decentralization is important, because centralized optimization problems may not be realistic to solve every time a user shows up on the site. In the remainder of the section we explain how a bidder may implement our Algorithm 1 via a form of decentralized double-pacing. We do this to explain it in generality, but we emphasize that the proxy-bidder setting is one of our main motivations.

Consider a bidder who attempts to maximize parity-regularized utility. By using Algorithm 1, the bidder will have two parameters

at time t : μ_t and λ_t , with some item of value v_t having appeared (the price p_t would not be known yet since the auction has not been run). Algorithm 1 requires that the bidder achieve an allocation that solves Equation (5), which we can rewrite as follows:

$$\begin{aligned} \arg \max_{x \in \{0,1\}} \{f_t(x) - \mu_t p_t x - \langle \lambda_t, c_t x \rangle\} \\ &= \arg \max_{x \in \{0,1\}} \{v_t x - (1 + \mu_t) p_t x - \langle \lambda_t, c_t \rangle x\} \\ &= \arg \max_{x \in \{0,1\}} \{(v_t - \langle \lambda_t, c_t \rangle) x - (1 + \mu_t) p_t x\} \\ &= \arg \max_{x \in \{0,1\}} \left\{ \frac{v_t - \langle \lambda_t, c_t \rangle}{1 + \mu_t} x - p_t x \right\}. \end{aligned}$$

From the above derivations, we observe at each t that there are three possible cases: (i) the bidder needs to win the item at time t if $\frac{v_t - \langle \lambda_t, c_t \rangle}{1 + \mu_t} > p_t$; (ii) the bidder has to pass on the item if the strict inequality is reversed; (iii) the bidder is indifferent if equality holds. Even though the bidder does not know p_t , this is exactly achieved within a second-price auction framework by bidding the value $\frac{v_t - \langle \lambda_t, c_t \rangle}{1 + \mu_t}$ at time t . Here, $1/(1 + \mu_t)$ plays the role of the standard multiplicative pacing parameter, but before multiplicative pacing, the bidder first applies an additive pacing term $\langle \lambda_t, c_t \rangle$.

Thus, we see that our approach can easily be adopted by individual bidders that wish to incorporate distributional regularization as part of their bidding procedure in a repeated-auction setting. For the same reason, proxy bidders can easily incorporate this in the case where an advertising platform wishes to enforce some amount of distributional parity while performing budget-pacing on behalf of the advertisers.

7 EXPERIMENTAL EVALUATION

In this section, we present numerical experiments on data from a large Internet advertising company.

Experimental setup. We construct a real-world dataset through logs of a large Internet advertising company. In this dataset, we have one advertiser (more specifically, we consider data at individual campaign level) participating in a sequence of auctions. We consider only ad requests coming from mobile devices. For each ad request we retrieve the advertiser's bid, the price computed via the auction mechanism, and the relevant categories for the request. We observe that, due to the nature of the bidding system, the bids provide a good approximation to the true advertiser's valuations. Bids and prices are suitably normalized for privacy reasons. Then, we generate a dataset consisting of 10^6 entries (i.e., one entry per ad requests) and, at each run of the algorithm, we sample from this dataset T requests without replacement. We consider two different ad-request breakdowns. The first one, denoted as breakdown B2, is a binary breakdown (i.e., $m = 2$) where the category is the *type* of mobile device originating the request (we consider the two most frequent such types). The second one, which we denote by B3, is a breakdown into three categories (i.e., $m = 3$) determined by the manufacturer of the mobile device. In all of the experiments we set ρ to be approximately equal to the 0.5 quantile of the empirical distribution of normalized prices from our dataset. In this way, the advertiser is set to win at most half of the ad requests which they are presented.

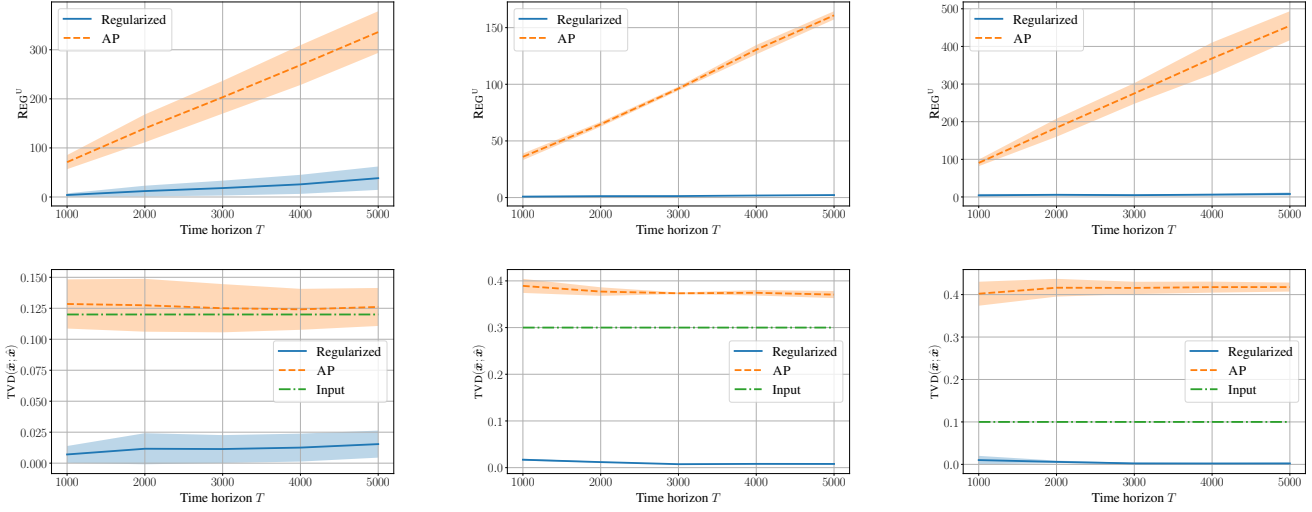


Figure 1: (Top-Left): Empirical upper bound on the regret (i.e., REG^U as defined in Equation (15)) for breakdown B2. (Bottom-Left): Total variation distance from the target distribution \hat{x} for B2. (Top-Center): Empirical upper bound on the regret for breakdown B3. (Bottom-Center): Total variation distance from the target distribution \hat{x} for B3. (Top-Right): Empirical upper bound on the regret for breakdown B2 with artificial prices unbalance. (Bottom-Right): Total variation distance from the target distribution \hat{x} for B2 with artificial prices unbalance. *Input* denotes the total variation distance between \hat{x} and the underlying empirical category distribution computed on the dataset.

Algorithms. We employ parity ray regularizers with $D_2(x; \bar{x}) = \|x - \bar{x}\|_2$ as the pseudo-distance measure. Moreover, given a sequence of input tuples $(v_t, p_t, c_t)_{t=1}^T$, let

$$\text{OPT}^U\left((v_t, p_t, c_t)_{t=1}^T\right) := \max_{x: x_t \in \mathcal{X}_t} \sum_{t=1}^T f_t(x_t) \text{ s.t. } \sum_{t=1}^T p_t x_t \leq T\rho$$

be the value of an optimal allocation in absence of regularization penalties. Then, we evaluate the regret guarantees of our algorithm by computing an upper bound to the parity-regularized regret $\text{OPT}(\mathcal{P}) - \text{REW}(\mathcal{P})$. By letting $s_T := (x_t, v_t, p_t, c_t)_{t=1}^T$ we compute

$$\text{REG}^U(s_T) := \text{OPT}^U\left((v_t, p_t, c_t)_{t=1}^T\right) - \text{REW}(s_T), \quad (15)$$

where $\text{REW}(s_T)$ is the realized reward for the sequence of online decisions x . The baseline algorithm which we employ is the *adaptive pacing algorithm* (AP) by Balseiro and Gur [14], which guarantees optimal regret in the stochastic setting without regularization. The step size is $\eta = 0.1/\sqrt{T}$. This choice guarantees the optimal order of regret both for our algorithm, and for adaptive pacing (in the non-regularized setting). Dual variables are initialized to 0. Convex programs are solved using the Xpress Optimization Suite 8.10 [27]. All experiments are run on a 24-core machine with 57 GB of RAM.

Results. We run 20 independent trials for each time horizon $T \in \{10^3, 2 \cdot 10^3, \dots, 5 \cdot 10^3\}$. The first row of Figure 1 reports the empirical upper bounds on the regret computed as in Equation (15). The second row of Figure 1 reports the *total variation distance* between the realized distributions of impressions \bar{x} and the target distribution \hat{x} . We compute this distance as $\text{TVD}(\bar{x}; \hat{x}) = \sup_j |\bar{x}_j - \hat{x}_j|$. The first column of Figure 1 reports result for the B2 breakdown. In this setting, we set $\hat{x} = (.5, .5)$, that is, the advertiser seeks to reach a uniform realized distribution of impressions over the two mobile device types. Algorithm 1 achieves small regret over all time

horizons, and guarantees a realized distribution of impression close to the target. On the other hand, the adaptive pacing algorithm shows linear regret in T , which is due to the penalty for not being able to steer the realized distribution of impression towards the target. The same behavior can be observed in the B3 setting (second column of Figure 1). Here, the empirical distribution of categories of the dataset is roughly $(.5, .3, .2)$. To make the problem harder, we try to skew the realized distribution of impressions towards the least frequent category by setting $\hat{x} = (.1, .3, .6)$. Algorithm 1 manages to keep the resulting distribution close to the target even in this challenging setting. It is interesting to notice that the distribution over categories of the input data would be closer to the target distribution than what achieved by the adaptive pacing algorithm. The same behavior can be observed when we artificially create an unbalance between the distribution of prices for the different categories. In these extreme cases, the adaptive pacing algorithm forces the advertiser to win only impression of the *cheap* category. This can be observed, for example, if we start from the B2 scenario described in Section 7, where we have two possible categories per item corresponding to different mobile device types, and then we artificially create an unbalance in the price distributions for the two categories, making one category slightly cheaper than the other. Figure 1 (Right-Column) reports the comparison between Algorithm 1 and AP. The total variation distance between the realized distribution of impressions \bar{x} and the target distribution \hat{x} is close to 0.5. Since $\hat{x} = (.5, .5)$ this means that AP is forcing the bidder to win *only* impressions for *cheap* requests. On the other hand, by employing Algorithm 1 we can effectively steer the realized distribution of impression towards the target without compromising the quality of the solution in terms of rewards.

REFERENCES

- [1] Deepak Agarwal, Souvik Ghosh, Kai Wei, and Siyu You. 2014. Budget pacing for targeted online advertisements at linkedin. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 1613–1619.
- [2] Shipra Agrawal and Nikhil R Devanur. 2014. Fast algorithms for online stochastic convex programming. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 1405–1424.
- [3] Shipra Agrawal, Nikhil R Devanur, and Lihong Li. 2016. An efficient algorithm for contextual bandits with knapsacks, and an extension to concave objectives. In *Conference on Learning Theory*. PMLR, 4–18.
- [4] Shipra Agrawal, Zizhuo Wang, and Yinyu Ye. 2014. A dynamic near-optimal algorithm for online linear programming. *Operations Research* 62, 4 (2014), 876–890.
- [5] Kareem Amin, Michael Kearns, Peter Key, and Anton Schwaighofer. 2012. Budget optimization for sponsored search: Censored learning in mdp. *arXiv preprint arXiv:1210.4847* (2012).
- [6] Alessandro Arlotto and Itai Gurvich. 2019. Uniformly bounded regret in the multisecretary problem. *Stochastic Systems* 9, 3 (2019), 231–260.
- [7] Vashist Avadhanula, Riccardo Colini-Baldeschi, Stefano Leonardi, Karthik Abinav Sankararaman, and Okke Schrijvers. 2021. Stochastic Bandits for Multi-platform Budget Optimization in Online Advertising. *arXiv preprint arXiv:2103.10246* (2021).
- [8] Ashwinkumar Badanidiyuru, John Langford, and Aleksandr Slivkins. 2014. Resourceful contextual bandits. In *Conference on Learning Theory*. PMLR, 1109–1134.
- [9] Santiago Balseiro, Anthony Kim, Mohammad Mahdian, and Vahab Mirrokni. 2017. Budget management strategies in repeated auctions. In *Proceedings of the 26th International Conference on World Wide Web*. 15–23.
- [10] Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. 2020. The Best of Many Worlds: Dual Mirror Descent for Online Allocation Problems. *arXiv preprint arXiv:2011.10124* (2020).
- [11] Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. 2020. Dual Mirror Descent for Online Allocation Problems. In *International Conference on Machine Learning*. PMLR, 613–628.
- [12] Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. 2020. Regularized Online Allocation Problems: Fairness and Beyond. *arXiv preprint arXiv:2007.00514* (2020).
- [13] Santiago R Balseiro, Omar Besbes, and Gabriel Y Weintraub. 2015. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science* 61, 4 (2015), 864–884.
- [14] Santiago R Balseiro and Yonatan Gur. 2019. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science* 65, 9 (2019), 3952–3968.
- [15] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Kamal Jain, Omid Etesami, and Mohammad Mahdian. 2007. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th international conference on World Wide Web*. 531–540.
- [16] Jonathan Borwein and Adrian S Lewis. 2010. *Convex analysis and nonlinear optimization: theory and examples*. Springer Science & Business Media.
- [17] Vincent Conitzer, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Eric Sodomka, Nicolas E Stier-Moses, and Chris Wilkens. 2019. Pacing Equilibrium in First-Price Auction Markets. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. 587–587.
- [18] Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E. Stier-Moses. 2021. Multiplicative Pacing Equilibria in Auction Markets. *Operations Research (to appear)* (2021). <https://arxiv.org/abs/1706.07151>
- [19] Nikhil R Devanur and Thomas P Hayes. 2009. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *Proceedings of the 10th ACM conference on Electronic commerce*. 71–78.
- [20] Nikhil R Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A Wilkens. 2019. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. *Journal of the ACM (JACM)* 66, 1 (2019), 1–41.
- [21] Nikhil R Devanur, Balasubramanian Sivan, and Yossi Azar. 2012. Asymptotically optimal algorithm for stochastic adwords. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. 388–404.
- [22] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review* 97, 1 (2007), 242–259.
- [23] eMarketer. 2020. eMarketer US Digital Ad Spending Update Q2 2020. <https://www.emarketer.com/content/us-digital-ad-spending-update-q2-2020>. Accessed: 2020-02-03.
- [24] Jon Feldman, Nitish Korula, Vahab Mirrokni, Shanmugavelayutham Muthukrishnan, and Martin Pál. 2009. Online ad assignment with free disposal. In *International workshop on internet and network economics*. Springer, 374–385.
- [25] Jon Feldman, Aranyak Mehta, Vahab Mirrokni, and Shan Muthukrishnan. 2009. Online stochastic matching: Beating 1-1/e. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, 117–126.
- [26] Jon Feldman, Shanmugavelayutham Muthukrishnan, Martin Pal, and Cliff Stein. 2007. Budget optimization in search-based advertising auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*. 40–49.
- [27] FICO. 2021. Xpress optimization suite. <https://www.fico.com/en/products/fico-xpress-optimization>
- [28] Lodewijk Gelauff, Ashish Goel, Kamesh Munagala, and Sravya Yandamuri. 2020. Advertising for Demographically Fair Outcomes. *arXiv preprint arXiv:2006.03983* (2020).
- [29] Gagan Goel and Aranyak Mehta. 2008. Online budgeted matching in random input models with applications to Adwords. In *SODA*, Vol. 8. 982–991.
- [30] Ramki Gummadi, Peter Key, and Alexandre Proutiere. 2013. Optimal bidding strategies and equilibria in dynamic auctions with budget constraints. *Available at SSRN 2066175* (2013).
- [31] Kartik Hosanagar and Vadim Cherepanov. 2008. Optimal bidding in stochastic budget constrained slot auctions. In *Proceedings of the 9th ACM Conference on Electronic Commerce*. 20–20.
- [32] Stefanus Jasin. 2015. Performance of an lp-based control for revenue management with unknown demand parameters. *Operations Research* 63, 4 (2015), 909–915.
- [33] Xiaocheng Li, Chunlin Sun, and Yinyu Ye. 2020. Simple and fast algorithm for binary integer and online linear programming. *arXiv preprint arXiv:2003.02513* (2020).
- [34] Xiaocheng Li and Yinyu Ye. 2019. Online linear programming: Dual convergence, new algorithms, and regret bounds. *arXiv preprint arXiv:1909.05499* (2019).
- [35] Mohammad Mahdian, Hamid Nazerzadeh, and Amin Saberi. 2012. Online optimization with uncertain information. *ACM Transactions on Algorithms (TALG)* 8, 1 (2012), 1–29.
- [36] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)* 54, 5 (2007), 22–es.
- [37] Vahab S Mirrokni, Shayan Oveis Gharan, and Morteza Zadimoghaddam. 2012. Simultaneous approximations for adversarial and stochastic online budgeted allocation. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*. SIAM, 1690–1701.
- [38] Milad Nasr and Michael Carl Tschantz. 2020. Bidding strategies with gender nondiscrimination constraints for online ad auctions. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 337–347.
- [39] Francesco Orabona. 2019. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213* (2019).
- [40] Hal R Varian. 2007. Position auctions. *International Journal of Industrial Organization* 25, 6 (2007), 1163–1178.
- [41] Di Wu, Xiujun Chen, Xun Yang, Hao Wang, Qing Tan, Xiaoxun Zhang, Jian Xu, and Kun Gai. 2018. Budget constrained bidding by model-free reinforcement learning in display advertising. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. 1443–1451.
- [42] Jian Xu, Kuang-chih Lee, Wentong Li, Hang Qi, and Quan Lu. 2015. Smart pacing for effective online ad campaign optimization. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2217–2226.
- [43] Weinan Zhang, Yifei Rong, Jun Wang, Tianchi Zhu, and Xiaofan Wang. 2016. Feedback control of real-time display advertising. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*. 407–416.
- [44] Weinan Zhang, Shuai Yuan, and Jun Wang. 2014. Optimal real-time bidding for display advertising. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 1077–1086.
- [45] Weinan Zhang, Ying Zhang, Bin Gao, Yong Yu, Xiaojie Yuan, and Tie-Yan Liu. 2012. Joint optimization of bid and budget allocation in sponsored search. In *Proceedings of the 18th ACM SIGKDD international conference on knowledge discovery and data mining*. 1177–1185.
- [46] Yunhong Zhou, Deeparnab Chakrabarty, and Rajan Lukose. 2008. Budget constrained bidding in keyword auctions and online knapsack problems. In *International Workshop on Internet and Network Economics*. Springer, 566–576.

A RELATED WORKS

Online allocation problems have been studied from a number of different perspectives. One stream of works addresses the case in which inputs are chosen by an adversary (see, e.g., Feldman et al. [24], Mehta et al. [36]). In our work, motivated by Internet advertising applications, we study the stochastic case in which inputs are drawn i.i.d. from an unknown distribution. Stochastic input models have been studied, among others, by Devanur and Hayes [19], Devanur et al. [21], Feldman et al. [25], Goel and Mehta [29], Mahdian et al. [35], Mirrokni et al. [37].

In the setting of stochastic input models with linear reward functions it is possible to guarantee an optimal order of regret of $O(T^{1/2})$ [4, 20]. More recent works by Li et al. [33] and Balseiro et al. [11] propose simple algorithms with $O(T^{1/2})$ regret guarantees that, in contrast with previous works, do not require periodically solving large linear programs. Finally, under some specific structural assumptions, it is possible to obtain a regret bound of order $O(\log T)$ [32, 34].

As discussed in Section 1 and Section 4, we study how to optimize for parity-regularized utilities, which are non-separable in nature. Two related works addressing similar problems are the following:

- Agrawal et al. [4] focus on how to solve general online stochastic convex programs that allow general concave objectives and convex constraints. However, their approach requires periodically solving convex optimization programs on historical data, unless the optimal value of the objective is known. Moreover, their algorithm handles resource constraints as *soft* constraints, while we are interested in enforcing a hard constraint on the budget.
- Balseiro et al. [12] propose a dual subgradient descent algorithm for a general class of non-separable objectives. In our work, we focus on the specific real-world problem of modeling advertisers' distributional preferences within standard pacing systems. Balseiro et al. [12] focus on a more general class of problems, without providing a concrete way to model our objective. In particular, we observe that the regularizer for *max-min fairness* which the authors propose in their Example 2 is not suitable for our setting. In order to employ their regularizer to model distributional preferences, an advertiser would have to appropriately pick parameters λ and ρ_i depending on the contingent state of the environment during the T time horizon. Specifically, this choice would require knowing beforehand the underlying input distribution of requests, as well as the *scale* of the number of ad opportunities that the advertiser will win during the T steps. On the other hand, the parity ray regularizer simplifies the interaction between the advertiser and the pacing mechanism, as the advertiser just needs to specify their target distribution, without requiring any additional knowledge on the current state of the environment. Moreover, the parity ray regularizer can be equipped with any pseudo-distance measure satisfying Assumption 2. Finally, in our paper we adapt the online mirror descent scheme by Balseiro et al. [10] to the case of non-separable objectives (which allows for a more general class of reference functions than the online gradient descent scheme of Balseiro et al. [12]). The

dual-descent scheme by Balseiro et al. [12] can be derived from our OMD scheme through an appropriate choice of the reference function. Our analysis draws heavily on the ideas developed in both Balseiro et al. [10] and Balseiro et al. [12]. By exploiting the specific structure of our problem we show how to relax some of the assumptions required by Balseiro et al. [10] (see, e.g., Assumption 2 of their paper).

There is also a stream of literature studying budget management problems within the framework of stochastic bandits with knapsacks constraints (see, e.g., the work by Avadhanula et al. [7]). The difference between our setting and the bandits literature is that, in online allocation problems, the advertiser observes the reward function and consumption matrix *before* making a decision, while in the bandit setting those are revealed *after* each decision is made. In order to deal with unknown reward functions and consumption matrices, the most frequent approach in contextual bandits is to discretize the context and action space. However, this inevitably leads to worse performance, and does not fit well with real-world pacing systems such as those described in Section 2. As an example, in Internet advertising applications bandits algorithms typically attain regret on the order of $O(T^{3/4})$ [3, 8].

Another line of research studies how individual bidders should optimize their budget spending across a set of auctions. This problem has been cast, for example, as a knapsack problem [15, 26, 46], a Markov Decision Process [5, 30, 41], a constrained optimization problem [44, 45], and an optimal control problem [42, 43].

In a recent work, Nasr and Tschantz [38] study how to compute bidding strategies for advertisers which aim at satisfying various types of parity constraints with respect to a binary user breakdown. Their model is different from ours since (i) we are interested in letting advertisers reach an arbitrary target distribution over an arbitrary user breakdown, while their goal is guaranteeing that specific parity constraints are satisfied for a binary user breakdown; (ii) we have a hard constraint on the budget, while they assume advertisers have an unlimited budget; (iii) the techniques which are adopted are very different since the authors model the problem as an MDP which is then solved with an ad-hoc method to properly manage the constraints, while we address distributional preferences directly within the pacing architecture.

B OMITTED PROOFS

LEMMA 5.1. *It holds*

$$\mathcal{D} := \{(\mu, \lambda) \in \mathbb{R}_{\geq 0} \times \mathbb{R}^m : -\infty < R^*(-\lambda) < +\infty\} = \mathbb{R}_{\geq 0} \times \mathbb{R}^m.$$

PROOF. Given $\lambda \in \mathbb{R}^m$, the following holds

$$\begin{aligned} R^*(-\lambda) &= \sup_{\bar{x} \in \Delta_+^m} \{R(\bar{x}) + \langle \lambda, \bar{x} \rangle\} \\ &\leq \sup_{\bar{x} \in \Delta_+^m} R(\bar{x}) + \sup_{\bar{x} \in \Delta_+^m} \langle \lambda, \bar{x} \rangle \\ &\leq \|\lambda\|_1 < +\infty, \end{aligned}$$

where the third inequality follows by Assumption 4.3. Analogously, given $\lambda \in \mathbb{R}^m$, we have

$$R^*(\lambda) \geq \underline{r} - \|\lambda\|_1 > -\infty,$$

which proves the statement. \square

THEOREM 5.2. *Given a time horizon $T \geq 0$, for any $(\mu, \lambda) \in \mathcal{D}$, $\text{OPT}(\mathcal{P}) \leq TD(\mu, \lambda|\mathcal{P})$.*

PROOF. Let $(v, p, c) = ((v_t)_{t=1}^T, (p_t)_{t=1}^T, (c_t)_{t=1}^T)$ be an arbitrary sequence of input tuples. From the definition of the baseline OPT we have that:

$$\begin{aligned} \text{OPT}(\mathcal{P}) &:= \mathbb{E}_{(v,p,c) \sim \mathcal{P}} \left[\begin{array}{l} \max_{x: x_t \in \mathcal{X}_t} \sum_{t=1}^T f_t(x_t) + TR \left(\frac{1}{T} \sum_{t=1}^T c_t x_t \right) \\ \text{s.t.} \quad \sum_{t=1}^T p_t x_t \leq T\rho \end{array} \right] \\ &= \mathbb{E}_{(v,p,c) \sim \mathcal{P}} \left[\begin{array}{l} \max_{\substack{x: x_t \in \mathcal{X}_t \\ \bar{x} \in \Delta_+^m}} \sum_{t=1}^T f_t(x_t) + TR(\bar{x}) \\ \text{s.t.} \quad \sum_{t=1}^T p_t x_t \leq T\rho \\ \bar{x} = \frac{1}{T} \sum_{t=1}^T c_t x_t \end{array} \right]. \end{aligned}$$

Then, by exploiting weak duality and the fact that input tuples are drawn i.i.d. from \mathcal{P} , we have that, for any $(\mu, \lambda) \in \mathcal{D}$,

$$\begin{aligned} \text{OPT}(\mathcal{P}) &\leq \mathbb{E}_{(v,p,c) \sim \mathcal{P}} \left[\max_{x, \bar{x}} \left\{ \sum_{t=1}^T f_t(x_t) + TR(\bar{x}) \right. \right. \\ &\quad \left. \left. - \mu \sum_{t=1}^T p_t x_t + T\mu\rho + \langle \lambda, T\bar{x} - \sum_{t=1}^T c_t x_t \rangle \right\} \right] \\ &\leq \mathbb{E}_{(v,p,c) \sim \mathcal{P}} \left[\sum_{t=1}^T f_t^*(\mu p_t + \langle \lambda, c_t \rangle) + TR^*(-\lambda) + T\rho\mu \right] \\ &= T \mathbb{E}_{(v_t, p_t, c_t) \sim \mathcal{P}} [f_t^*(\mu p_t + \langle \lambda, c_t \rangle) + R^*(-\lambda)] + \rho\mu \\ &= TD(\mu, \lambda|\mathcal{P}). \end{aligned}$$

This proves our statement. \square

LEMMA 5.3. *Consider an arbitrary $\mathcal{P} \in \Delta^T$. For each $t \in [T]$, let x_t and \bar{x}_t be computed according to Equation (5) and Equation (6), respectively. Moreover, given the time of budget depletion τ , let $\bar{\mu} = 1/\tau \sum_{t=1}^{\tau} \mu_t$ and $\bar{\lambda} = 1/\tau \sum_{t=1}^{\tau} \lambda_t$. Then, it holds*

$$\begin{aligned} \mathbb{E}_{\mathcal{P}} \left[\sum_{t=1}^{\tau} (f_t(x_t) + R(\bar{x}_t)) \right] &\geq \\ \mathbb{E}_{\mathcal{P}} \left[\tau D(\bar{\mu}, \bar{\lambda}|\mathcal{P}) - \sum_{t=1}^{\tau} (\mu_t(\rho - p_t x_t) - \langle \lambda_t, \bar{x}_t - c_t x_t \rangle) \right]. \end{aligned}$$

PROOF. For each $t \leq \tau$, by Equation (5) and Equation (6) we have

$$\begin{aligned} &\mathbb{E}_{\mathcal{P}} [f_t(x_t) + R(\bar{x}_t)] \\ &= \mathbb{E} [f_t^*(\mu_t p_t + \langle \lambda_t, c_t \rangle) + \mu_t p_t x_t + \langle \lambda_t, c_t x_t \rangle + R^*(\lambda_t) - \langle \lambda_t, \bar{x}_t \rangle] \\ &= \mathbb{E} [D(\mu_t, \lambda_t|\mathcal{P}) - \rho\mu_t + \mu_t p_t x_t + \langle \lambda_t, c_t x_t \rangle - \langle \lambda_t, \bar{x}_t \rangle] \\ &= \mathbb{E} [D(\mu_t, \lambda_t|\mathcal{P}) - \mu_t(\rho - p_t x_t) - \langle \lambda_t, \bar{x}_t - c_t x_t \rangle], \end{aligned}$$

where the second equality holds by Equation (4). By summing over $t = 1, \dots, \tau$ we obtain

$$\begin{aligned} \mathbb{E}_{\mathcal{P}} \left[\sum_{t=1}^{\tau} f_t(x_t) + R(\bar{x}_t) \right] &= \sum_{t=1}^{\tau} \mathbb{E} [f_t(x_t) + R(\bar{x}_t)] \\ &= \mathbb{E} \left[\sum_{t=1}^{\tau} (D(\mu_t, \lambda_t|\mathcal{P}) - \mu_t(\rho - p_t x_t) - \langle \lambda_t, \bar{x}_t - c_t x_t \rangle) \right] \\ &\geq \mathbb{E} \left[\tau D(\bar{\mu}, \bar{\lambda}|\mathcal{P}) - \sum_{t=1}^{\tau} (\mu_t(\rho - p_t x_t) - \langle \lambda_t, \bar{x}_t - c_t x_t \rangle) \right], \end{aligned}$$

where the last inequality holds by convexity of the Lagrangian dual function. Specifically, $f_t^*(\mu p + \langle \lambda, c \rangle)$ is the composition of a convex function (by definition of conjugate function) with an affine mapping, which is convex. Therefore, D is convex since it is the result of a sum of convex functions, and expectation preserves convexity. \square