

Learning in Repeated Auctions with Budgets: Regret Minimization and Equilibrium

Santiago R. Balseiro
Columbia University

Yonatan Gur*
Stanford University

June 28, 2018

Abstract

In online advertising markets, advertisers often purchase ad placements through bidding in repeated auctions based on realized viewer information. We study how budget-constrained advertisers may compete in such sequential auctions in the presence of uncertainty about future bidding opportunities and competition. We formulate this problem as a sequential game of incomplete information, where bidders know neither their own valuation distribution, nor the budgets and valuation distributions of their competitors. We introduce a family of practical bidding strategies we refer to as *adaptive pacing* strategies, in which advertisers adjust their bids according to the sample path of expenditures they exhibit, and analyze the performance of these strategies in different competitive settings. We establish the asymptotic optimality of these strategies when competitors' bids are independent and identically distributed over auctions, but also when competing bids are arbitrary. When all the bidders adopt these strategies, we establish the convergence of the induced dynamics and characterize a regime (well motivated in the context of online advertising markets) under which these strategies constitute an approximate Nash equilibrium in dynamic strategies: the benefit from unilaterally deviating to other strategies, including ones with access to complete information, becomes negligible as the number of auctions and competitors grows large. This establishes a connection between regret minimization and market stability, by which advertisers can essentially follow approximate equilibrium bidding strategies that also ensure the best performance that can be guaranteed off equilibrium.

KEYWORDS: Sequential auctions, online advertising, online learning, stochastic optimization, stochastic approximation, incomplete information, regret analysis, dynamic games

1 Introduction

Online ad spending has grown dramatically in recent years, reaching over \$70 billion in the United States in 2016 (eMarketer 2016b). A substantial part of this spending is generated through auction platforms where advertisers sequentially submit bids to place their ads at slots as these become available. Important examples include ad exchange platforms such as Google's DoubleClick and

*The authors are grateful to Omar Besbes, Negin Golrezaei, Vahab S. Mirrokni, Michael Ostrovsky, and Gabriel Y. Weintraub for their valuable comments. Correspondence: srb2155@columbia.edu, ygur@stanford.edu.

Yahoo’s RightMedia, more centralized platforms such as Facebook Exchange and Twitter Ads, as well as sponsored search platforms such as Google’s AdWords and Microsoft’s Bing Ads. In such markets, advertisers may participate in many thousands of auctions per day, where competitive interactions are *interconnected* by budgets that limit the total expenditures of advertisers throughout their campaign. The complexity of this competition, together with the frequency of opportunities and the time scale on which decisions are made, are key drivers in the rise of automated bidding algorithms that have become common practice among advertisers (eMarketer 2016a).

One *operational* challenge that is fundamental for managing budgeted campaigns in online ad markets is to balance present and future bidding opportunities effectively. For example, advertisers face the analytical and computational challenge of timing the depletion of their budget with the planned termination of their campaign since running out of budget prior to the campaign’s planned termination (or alternatively, reaching the end of the campaign with unutilized funds) may result in significant opportunity loss. The challenge of accounting for future opportunities throughout the bidding process has received increasing attention in the online advertising industry, as platforms now recommend advertisers to “pace” the rate at which they spend their budget, and provide *budget-pacing* services that are based on a variety of heuristics.¹

Budget pacing is particularly challenging due to the many uncertainties that characterize the competitive landscape in online ad markets. First, while advertisers can evaluate, just before each auction, the value of the *current* opportunity (in terms of, e.g., likelihood of purchase), advertisers do not hold such information about *future* opportunities that may be realized throughout the campaign. Second, advertisers typically know very little about the extent of the competition they face in online ad markets. They typically do not know the number of advertisers that will bid in each auction, nor important characteristics of their competitors such as their budgets and their expectations regarding future opportunities. In addition, there is uncertainty about the varying levels of competitors’ strategic and technical sophistication: while some firms devote considerable resources to develop complex algorithms that are designed to identify and respond to strategies of competitors in real time to maximize their campaign utility, other, less resourceful firms may adopt simple bidding strategies that may be independent of idiosyncratic user information, and that may even appear arbitrary.

¹A descriptive discussion of the budget pacing service Facebook offers advertisers appears at <https://developers.facebook.com/docs/marketing-api/pacing>; See also recommendations for advertisers in online blogs powered by Twitter Ads (<https://dev.twitter.com/ads/campaigns/budget-pacing>), Google’s ad exchange platform DoubleClick (<https://support.google.com/ds/answer/2682462>), as well as ExactDrive (<http://exactdrive.com/news/how-budget-pacing-helps-your-online-advertising-campaign>).

The main questions we address in this paper are: (i) How should budget-constrained advertisers compete in repeated auctions under uncertainty? (ii) What type of performance can be guaranteed *with* and *without* typical assumptions on competitors’ behavior? (iii) Can a single class of strategies constitute an equilibrium (when adopted by all advertisers) while achieving the best performance that can be guaranteed off equilibrium? These questions are of practical importance and draw attention from online advertisers and ad platforms.

1.1 Main contribution

The main contribution of the paper lies in introducing a new family of practical and intuitive bidding strategies that dynamically adapt to uncertainties and competition throughout an advertising campaign based only on observed expenditures, and in analyzing the performance of these strategies *off equilibrium* and *in equilibrium*. In more detail, our contribution is along the following lines.

Formulation and solution concepts. We formulate the budget-constrained competition between advertisers as a sequential game of incomplete information, where advertisers know neither their own valuation distribution, nor the budgets and valuation distributions of their competitors. To evaluate the performance of a given strategy without any assumptions on competitors’ behavior, we quantify the portion it can guarantee of the performance of the best (feasible) dynamic sequence of bids one could have selected with the benefit of hindsight. The proposed performance metric, referred to as γ -competitiveness, extends that of *no-regret* strategies (commonly used in *unconstrained* settings) to a more stringent benchmark that allows for dynamic sequences of actions subject to global constraints. By showing that the maximal portion that can be guaranteed is, in general, less than one, we establish the impossibility of a “no-regret” notion relative to a dynamic benchmark.

Adaptive pacing strategies. We introduce a new class of practical bidding strategies that adapt the bidding behavior of advertisers throughout the campaign, based only on the sample path of expenditures they exhibit. We refer to this class as *adaptive pacing* strategies, as these strategies dynamically adjust the “pace” at which the advertiser spends its budget. This pace is controlled by a single parameter: a dual multiplier that determines the extent to which the advertiser bids below its true values. This multiplier is updated according to an intuitive primal-dual scheme that extends single-agent learning ideas for the purpose of balancing present and future opportunities in uncertain environments. In particular, these strategies aim at maximizing instantaneous payoffs while guaranteeing that the advertiser depletes its budget “close” to the end of the campaign.

Performance and stability. We analyze the performance of adaptive pacing strategies under different assumptions on the competitive environment. On the one hand, when competitors’ bids are independent and identically distributed over auctions (we refer to this setting as *stationary* competition), the performance of this class of strategies converges, in a long-run average sense, to the best performance attainable with the benefit of hindsight. On the other hand, when competitors’ bids are arbitrary (we refer to this setting as *arbitrary* competition), we establish, through matching lower and upper bounds on γ -competitiveness, the asymptotic optimality of these strategies; i.e., no other strategy can guarantee a larger portion of the best performance in hindsight. Together, these results establish performance guarantees and asymptotic optimality in both the “optimistic” and “pessimistic” scenarios. When all bidders follow adaptive pacing strategies competing bids are non-stationary and endogenous. We show that in such a scenario the sequences of dual multipliers (and the resulting payoffs) converge, and characterize a regime (that is well grounded in practice) under which these strategies constitute an approximate Nash equilibrium: the benefit from unilaterally deviating to any *fully informed* dynamic strategy diminishes to zero as the number of auctions and the market size grow large, even when advertisers know their value distribution as well as the value distributions and budgets of their competitors up front, and even when they can acquire real-time information on the past value realizations, bids, and payoffs of their competitors.

Notably, without using any prior knowledge of the competitive setting (that is, whether it is stationary, arbitrary, or of simultaneous learning), these strategies asymptotically attain the best performance that is achievable when advertisers have ex ante knowledge of the nature of the competition. Our results (summarized in Table 1) establish a connection between performance and market stability in our setting: advertisers can follow bidding strategies that constitute an approximate equilibrium and that, at the same time, asymptotically ensure the best performance that can be guaranteed off equilibrium.

Nature of competition		
Stationary	Arbitrary	Simultaneously learning
Long-run average convergence to the best performance in hindsight	Long-run average convergence to the largest attainable portion of the best performance in hindsight	Convergence in multipliers and in performance; approximate equilibrium in large markets

Table 1: Performance characteristics of adaptive pacing strategies

To assess the value that may be captured in practice by adaptive pacing strategies, we show empirically, using data from a large online ad auction platform, that these strategies often achieve performance that is very close to the best performance that could have been achieved in hindsight.

1.2 Related work

Our study lies in the intersection of marketing, economics, operations, and computer science. While we focus on the broad application area of online advertising, our analysis includes methodological aspects of learning, games of incomplete information, and auctions theory.

Sequential games of incomplete information. There is a rich literature stream on the conditions under which observation-based dynamics may or may not converge to an equilibrium. For an extensive overview of related work, which excludes budget considerations and mostly focuses on discrete action sets and complete information, see the books of Fudenberg and Levine (1998) and Hart and Mas-Colell (2013). While convergence to equilibrium has appealing features, it generally does not imply any guarantees on performance along the strategy’s decision path. One important notion that has been used to evaluate performance along the decision path in many sequential settings is Hannan consistency (Hannan 1957), also referred to as universal consistency and, equivalently, no-regret. A strategy is Hannan consistent if, when competitors’ actions are arbitrary, it guarantees a payoff at least as high as any static action that could have been selected in hindsight. An important connection between Hannan consistency and convergence to equilibrium is established by Hart and Mas-Colell (2000), who provide a no-regret strategy that, when adopted by all players, induces an empirical distribution of play that converges to a correlated equilibrium.

In the present paper we advance similar ideas by establishing a connection between performance and equilibrium in a setting with *global constraints* that limit the sequence of actions players may take. Since in such a setting repeating any single action may be infeasible as well as generate poor performance, we replace the *static* benchmark that appears in Hannan consistency with the best *dynamic* and *feasible* sequence of actions. We characterize the maximal portion of this dynamic benchmark that may be guaranteed, provide a class of strategies that achieve this portion, and demonstrate practical settings in which these strategies constitute an ε -Nash equilibrium among fully informed dynamic strategies. This equilibrium notion is particularly ambitious under incomplete information as it involves look-ahead considerations, and is hard to maintain even in more information-abundant settings (see, e.g., Brafman and Tennenholtz 2004 and Ashlagi et al. 2012).

Learning and equilibrium in sequential auctions. Our work relates to other papers that study learning in various settings of repeated auctions. Iyer et al. (2014) adopt a mean-field approximation to study repeated auctions in which bidders learn about their own private value over time. Weed et al. (2016) consider a sequential auction setting, where bidders do not know

the associated intrinsic value, and adopt multi-armed bandit policies to analyze the underlying exploration-exploitation tradeoff. Bayesian learning approaches are analyzed, e.g., by Hon-Snir et al. (1998) in the context of sequential first-price auctions, as well as by Han et al. (2011) under monitoring and entry costs. These studies address settings significantly different from ours and, in particular, do not consider budget constraints. In our setting, items are sold via repeated second-price auctions and bidders observe current private values when bidding (but are uncertain about future values). While truthful bidding is weakly dominant in a static second-price auction, a budget-constrained bidder needs to learn and account for the (endogenous) option value of future auctions, which depends directly on its budget, together with unknown factors such as its value distribution as well as the budgets, value distributions, and strategic behavior of its competitors.

Repeated auctions with budgets. Balseiro et al. (2015) introduce a fluid mean-field approximation to study the outcome of the strategic interaction between budget-constrained advertisers in a complete information setting. They show that stationary strategies that shade bids using a constant multiplier can constitute an approximate Nash equilibrium, but these strategies rely on the value distributions and the budget constraints of *all* the advertisers. The present paper focuses on the practical challenge of bidding in these repeated auctions and, in particular, when such broad information on the entire market is not available to the bidder, and when competitors do not necessarily follow an equilibrium strategy. While we do not impose a fluid mean-field approximation in the present paper, under such an approximation our proposed strategies could be viewed as converging under *incomplete* information to an equilibrium in fluid-based strategies. In a related work, Conitzer et al. (2017) introduce a pacing equilibrium concept that also rely on complete information. They provide evidence through numerical experiments that the adaptive pacing strategies introduced here would converge to an equilibrium in their setting as well. Convergence under realistic (incomplete) information of the simple dynamics introduced here to the different notions of equilibria mentioned above provides some support to these equilibria as solution concepts. A few other budget management approaches have been recently considered in the literature; see, e.g., Karande et al. (2013), Charles et al. (2013), and Balseiro et al. (2017).

Zhou et al. (2008) study budget-constrained bidding for sponsored search in an adversarial setting and provide an algorithm with a competitive ratio depending on upper and lower bounds on the value-to-weight ratio. Our results are not directly comparable with theirs. Considering a stationary setting, Badanidiyuru et al. (2013) study multi-armed bandit problems in which the decision maker may be resource-constrained and discuss applications to bidding in repeated auctions with budgets

where arms correspond to possible bids. In our paper the information structure is different since advertisers observe their values before bidding. Our paper also relates to the literature on contextual multi-armed bandit problems with resource constraints, where the context corresponds to the impression value (see, e.g., Badanidiyuru et al. 2014). The algorithm of Badanidiyuru et al. (2014) can be applied to our setting by discretizing the action space, with performance guarantees that are worse than the guarantees obtained in the present paper. Borgs et al. (2007) study heuristics for bidding in repeated auctions in which advertisers simultaneously update bids based on past campaigns and establish convergence across campaigns. By contrast, we study strategies that converge throughout a single campaign. Finally, Jiang et al. (2014) analyze numerically an approximation scheme that is similar to ours, but from the perspective of a single agent that bids in the presence of an exogenous and stationary market. To the best of our knowledge, the present paper is the first to propose a class of practical budget-pacing strategies with performance guarantees in stationary and adversarial settings, and concrete notions of optimality under competition.

Stochastic and convex optimization. The bidding strategies we suggest are based on approximating a Lagrangian dual objective by constructing and following sub-gradient estimates in the dual space. In this respect our strategies relate to the class of mirror descent schemes introduced by Nemirovski and Yudin (1983); see also Beck and Teboulle (2003) and Chapter 2 of Shalev-Shwartz (2012), as well as the potential-based gradient descent method in Chapter 11 of Cesa-Bianchi and Lugosi (2006). More broadly, these strategies belong to the class of stochastic approximation methods that have been widely studied and applied in a variety of areas; for a comprehensive review see Benveniste et al. (1990), Kushner and Yin (2003), and Lai (2003), as well as Araman and Caldentey (2011) for more recent applications in revenue management. This strand of the literature focuses on single-agent learning methods for adjusting to uncertain environments that are *exogenous* and *stationary*. Exceptions that consider the performance of stochastic approximation schemes in non-stationary environments include the formulations in §3.2 of Kushner and Yin (2003), Chapter 4 of Benveniste et al. (1990), as well as Besbes et al. (2015). Notably, these frameworks consider underlying environments that change *exogenously*, and decision makers that operate as monopolists. In other related works, Rosen (1965) demonstrates finding an equilibrium in concave games by applying a gradient-based method in a centralized manner with complete information on the payoff functions, and Nedic and Ozdaglar (2009) study the convergence of distributed gradient-based methods adopted by cooperative agents.

The present paper contributes to this literature by extending, in two key aspects, single-agent

learning ideas in a setting of practical importance. First, our proposed strategies use expenditure observations to construct sub-gradient estimates period by period, each time for a *different component* of the dual objective we approximate. This facilitates efficient learning *throughout* the campaign, rather than between campaigns. Second, the proposed strategies are analyzed in settings that include both uncertainty and competition. In particular, when these strategies are adopted by all advertisers, the resulting environment is both endogenous and non-stationary.

Our performance analysis in the presence of arbitrary competition relates to the literature initiated by Zinkevich (2003) on Online Convex Optimization in an *adversarial* framework, where underlying cost functions can be selected at any point in time by an adversary, depending on the actions of the decision maker. (This constitutes a more pessimistic environment than the traditional stochastic setting, where a cost function is picked a priori and held fixed.) For an overview of related settings see Chapter 2 of Shalev-Shwartz (2012).

Adaptive algorithms. The challenge of designing strategies that adapt to unknown characteristics of the environment (in the sense of achieving ex post performance that is essentially as good as the performance that can be achieved when advertisers have ex ante knowledge of the nature of the competition) dates back to studies in the statistics literature (see Tsybakov 2008 and references therein), and has seen recent interest in machine learning. Examples include Mirrokni et al. (2012) who seek to design an algorithm for an online budgeted allocation problem that achieves an optimal competitive ratio in both adversarial and stochastic settings, Seldin and Slivkins (2014) who present an algorithm that achieves (near) optimal performance in both stochastic and adversarial multi-armed bandit regimes without prior knowledge of the nature of the environment, and Sani et al. (2014) who consider an online convex optimization setting and derive algorithms that are rate optimal regardless of whether the target function is weakly or strongly convex.

2 Incomplete information model

We model a sequential game of incomplete information with budget constraints, where K risk-neutral heterogeneous advertisers repeatedly bid to place ads through second-price auctions. We note that many of our modeling assumptions can be generalized and are made only to simplify exposition and analysis; some generalizations are discussed in §2.1.

In each time period $t = 1, \dots, T$ there is an available ad slot that is auctioned. The auctioneer (i.e., the platform) first shares with the advertisers some information about a visiting user, and

then runs a second-price auction with a reserve price of zero to determine which ad to show to the user. The information provided by the auctioneer heterogeneously affects the value advertisers perceive for the impression based on their targeting criteria. The values advertisers assign to the impression they bid for at time t are denoted by the random vector $(v_{k,t})_{k=1}^K$ and are assumed to be independently distributed across impressions and advertisers, with a support over $\prod_{k=1}^K [0, \bar{v}_k] \subset \mathbb{R}_+^K$. We denote the marginal cumulative distribution function of $v_{k,t}$ by F_k , and assume that F_k is absolutely continuous with bounded density f_k . Advertiser k has a budget B_k , which limits the total payments that can be made by the advertiser throughout the campaign. We denote by $\rho_k := B_k/T$ the target expenditure (per impression) rate of advertiser k . We assume that $0 < \rho_k \leq \bar{v}_k$ for each advertiser $k \in \{1, \dots, K\}$; otherwise the problem would be trivial as the advertiser would never deplete its budget. Each advertiser k is therefore characterized by a type $\theta_k := (F_k, \rho_k)$.

In each period t each advertiser privately observes a valuation $v_{k,t}$, and then posts a bid $b_{k,t}$. Given the bidding profile $(b_{k,t})_{k=1}^K \in \mathbb{R}_+^K$, we denote by $d_{k,t}$ the highest competing bid faced by advertiser k :

$$d_{k,t} := \max_{i: i \neq k} \{b_{i,t}\}.$$

We assume that advertisers have a quasilinear utility function given by the difference between the sum of the valuations generated by the impressions won and the expenditures corresponding to the second-price rule. We denote by $z_{k,t}$ the expenditure of advertiser k at time t :

$$z_{k,t} := \mathbf{1}\{d_{k,t} \leq b_{k,t}\}d_{k,t},$$

and by $u_{k,t} := \mathbf{1}\{d_{k,t} \leq b_{k,t}\}(v_{k,t} - d_{k,t})$ the corresponding net utility.² After the bidding takes place, the auctioneer allocates the ad to the highest bidder, and each advertiser k privately observes its own expenditure $z_{k,t}$ and net utility $u_{k,t}$.

Information structure and admissible budget-feasible bidding strategies. We assume that at the beginning of the horizon advertisers have no information on their valuation distributions, nor on their competitors'; that is, the distributions $\{F_k : k = 1 \dots, K\}$ are unknown. Each advertiser knows its own target expenditure rate ρ_k , as well as the length of the campaign T , but has no information on the budgets or the target expenditure rates of its competitors. In particular, each advertiser does not know the number of competitors in the market, their types, or its own type.

²Our results hold for random or lexicographic tie-breaking rules. To simplify exposition we assume that ties are broken in favor of the decision maker.

We next formalize the class of (potentially randomized) admissible and budget-feasible bidding strategies. Let y be a random variable defined over a probability space $(\mathbb{Y}, \mathcal{Y}, \mathbb{P}_y)$. We denote by $\mathcal{H}_{k,t}$ the *history* available at time t to advertiser k , defined by

$$\mathcal{H}_{k,t} := \sigma \left(\left(v_{k,\tau}, b_{k,\tau}, z_{k,\tau}, u_{k,\tau} \right)_{\tau=1}^{t-1}, v_{k,t}, y \right),$$

for any $t \geq 2$, with $\mathcal{H}_{k,1} := \sigma(v_{k,1}, y)$. A bidding strategy for advertiser k is a sequence of mappings $\beta = (\beta_1, \beta_2, \dots)$, where the functions $\beta_1 : \mathbb{R}_+ \times \mathbb{Y} \rightarrow \mathbb{R}_+$ and $\beta_t : \mathbb{R}_+^{4(t-1)+1} \times \mathbb{Y} \rightarrow \mathbb{R}_+$ for $t = 2, \dots, T$ map available histories to bids $\{b_{k,t}^\beta : t = 1, \dots, T\}$. A strategy is *admissible*, or *non-anticipating*, if it depends only on information that is available to the bidder; that is, if in each period t the strategy β_t is measurable with respect to the filtration $\mathcal{H}_{k,t}$. We say that β is *budget feasible* if it generates expenditures that are constrained by the available budget; that is, $\sum_{t=1}^T \mathbf{1}\{d_{k,t} \leq b_{k,t}^\beta\} d_{k,t} \leq \rho_k T$ for *any* realized vector of highest competitors' bids $\mathbf{d}_k = (d_{k,t})_{t=1}^T \in \mathbb{R}_+^T$. We denote by \mathcal{B}_k the class of admissible budget-feasible strategies for advertiser k with target expenditure rate ρ_k , and by $\mathcal{B} = \mathcal{B}_1 \times \dots \times \mathcal{B}_K$ the product space of strategies for all advertisers.

Off-equilibrium performance. To evaluate the off-equilibrium performance of a candidate strategy $\beta \in \mathcal{B}_k$ without any assumptions on the competitors' strategies, we define the sample-path performance *given* vectors of realized values $\mathbf{v}_k = (v_{k,t})_{t=1}^T \in [0, \bar{v}_k]^T$ and realized competing bids \mathbf{d}_k as follows:

$$\pi_k^\beta(\mathbf{v}_k; \mathbf{d}_k) := \mathbb{E}^\beta \left[\sum_{t=1}^T \mathbf{1}\{d_{k,t} \leq b_{k,t}^\beta\} (v_{k,t} - d_{k,t}) \right],$$

where the expectation is taken with respect to any randomness embedded in the strategy.³

We quantify the minimal loss a strategy can *guarantee* relative to any sequence of bids one could have made with the benefit of *hindsight*. Given sequences of realized valuations \mathbf{v}_k and highest competing bids \mathbf{d}_k , we denote by $\pi_k^H(\mathbf{v}_k; \mathbf{d}_k)$ the best performance advertiser k could have achieved with the benefit of hindsight:

$$\begin{aligned} \pi_k^H(\mathbf{v}_k; \mathbf{d}_k) := \max_{\mathbf{x}_k \in \{0,1\}^T} & \sum_{t=1}^T x_{k,t} (v_{k,t} - d_{k,t}) \\ \text{s.t.} & \sum_{t=1}^T x_{k,t} d_{k,t} \leq T \rho_k, \end{aligned} \tag{1}$$

³While the excess budget remaining after period T is not included in this payoff, it can be accounted for without impacting our results.

where the binary variable $x_{k,t}$ indicates whether the auction at time t is won by advertiser k . The solution of (1) is the best *dynamic* response (sequence of bids that could have been made) against a fixed environment characterized by \mathbf{v}_k and \mathbf{d}_k . Notably, if \mathbf{v}_k and \mathbf{d}_k are such that the bidder is *effectively* unconstrained (that is, if $\sum_{t=1}^T \mathbf{1}\{d_{k,t} \leq v_{k,t}\} d_{k,t} \leq \rho_k T$), then π_k^H is achieved by truthful bidding (i.e., $b_{k,t} = v_{k,t}$); otherwise, π_k^H is the solution to a knapsack problem (for an overview of knapsack problems see, e.g., Martello and Toth 1990). While the identification of the best sequence of bids in hindsight does not account for the potential competitive reaction to that sequence, this benchmark is tractable and has appealing practical features. For example, the best performance in hindsight can be computed using only historical data and requires no behavioral assumptions on competitors (see, e.g., Talluri and van Ryzin 2004, Ch. 11.3).

Non-anticipating strategies might not be able to achieve or approach π_k^H . For some $\gamma \in [1, \infty)$, a bidding strategy $\beta \in \mathcal{B}_k$ is said to be asymptotic γ -competitive (Borodin and El-Yaniv 1998) if

$$\limsup_{\substack{T \rightarrow \infty, \\ B_k = \rho_k T}} \sup_{\substack{\mathbf{v}_k \in [0, \bar{v}_k]^T \\ \mathbf{d}_k \in \mathbb{R}_+^T}} \frac{1}{T} \left(\pi_k^H(\mathbf{v}_k; \mathbf{d}_k) - \gamma \pi_k^\beta(\mathbf{v}_k; \mathbf{d}_k) \right) \leq 0. \quad (2)$$

An asymptotic γ -competitive bidding strategy (asymptotically) guarantees a portion of at least $1/\gamma$ of the performance of any dynamic sequence of bids that could have been selected in hindsight.⁴

Equilibrium solution concept. Given a strategy profile $\beta = (\beta_k)_{k=1}^K \in \mathcal{B}$, we denote by Π_k^β the total expected payoff for advertiser k :

$$\Pi_k^\beta := \mathbb{E}_{\mathbf{v}}^\beta \left[\sum_{t=1}^T \mathbf{1}\{d_{k,t}^{\beta_{-k}} \leq b_{k,t}^{\beta_k}\} (v_{k,t} - d_{k,t}^{\beta_{-k}}) \right],$$

where the expectation is taken with respect to any randomness embedded in the strategies β and the random values of all the advertisers, and where we denote the highest competing bid induced by competitors' strategies $\beta_{-k} = (\beta_i)_{i \neq k}$ by $d_{k,t}^{\beta_{-k}} = \max_{i: i \neq k} \{b_{i,t}^{\beta_i}\}$. We say that a strategy profile $\beta \in \mathcal{B}$ constitutes an ε -Nash equilibrium in dynamic strategies if each player's incentive to unilaterally deviate to another strategy is at most ε , that is, if for all k :

⁴While in the definition of $\pi_k^\beta(\mathbf{v}_k; \mathbf{d}_k)$, the vectors \mathbf{v}_k in $[0, \bar{v}_k]^T$ and \mathbf{d}_k in \mathbb{R}_+^T are fixed in advance, we also allow the components of these vectors to be selected dynamically according to the realized path of β ; that is, $v_{k,t}$ and $d_{k,t}$ could be measurable with respect to the filtration $\sigma(\{v_{k,\tau}, b_{k,\tau}, z_{k,\tau}, u_{k,\tau}\}_{\tau=1}^{t-1})$. This allows nonoblivious or adaptive adversaries and, in particular, valuations and competing bids that are affected by the player's strategy. For further details see Chapter 7 of Cesa-Bianchi and Lugosi (2006).

$$\frac{1}{T} \left(\sup_{\tilde{\beta} \in \mathcal{B}_k^{\text{CI}}} \Pi_k^{\tilde{\beta}, \beta_k} - \Pi_k^{\beta} \right) \leq \varepsilon, \quad (3)$$

where $\mathcal{B}_k^{\text{CI}} \supseteq \mathcal{B}_k$ denotes the class of strategies with *complete information* on market primitives. While a more precise description of this class will be provided in §5, we note that this class also includes strategies with access to complete information on the types $(\theta_i)_{i=1}^K$ and on the bids made in past periods by all advertisers. Naturally, the solution concept in (3) takes into account the strategic response of the competitive environment to the actions of an advertiser.

2.1 Discussion of model assumptions

Although we focus on second-price auctions, our results hold for all deterministic dominant-strategy mechanisms that are incentive compatible and individually rational with non-negative transfers. This broad class includes mechanisms such as second-price auctions with anonymous reserve prices, second-price auctions with personalized reserve prices, ranking mechanisms, and Myerson’s optimal mechanism. A question of theoretical and practical interest is how to extend the adaptive pacing approach that is suggested in this paper to non-truthful mechanisms such as first-price auctions.

Predicated on the common practice of running campaigns with daily budgets, we focus on *synchronous* campaigns that start and finish simultaneously. The synchronicity of the campaigns allows one to capture some key market features, while providing the tractability required to transparently highlight key problem characteristics. We note that the results in §3.3 that are established under arbitrary competition are agnostic to this assumption and hold even when campaigns start and finish at different times. An interesting research direction is to extend our convergence and equilibrium results to models with asynchronous campaigns such as those in Iyer et al. (2014) and Balseiro et al. (2015).

Our model assumes that values are independent across time and advertisers, and absolutely continuous. The assumption that values are independent across time is frequently used in the literature and is based on typical browsing behavior of internet users that leads to weak inter-temporal correlations. Our model can be relaxed to accommodate correlation across advertisers’ values as long as certain technical conditions (such as smoothness of the dual objective function and strong monotonicity of the expenditure function) hold. Although we provide sufficient conditions that rely on values being independent, these conditions can be shown to hold when values are imperfectly correlated, and when distributions have atoms at zero (capturing the possibility of advertisers not participating in every auction). We discuss the absolute continuity of values in §6.

In our model the number of impressions is known to all advertisers. In practice advertisers may infer the length of their campaign (in terms of total number of impressions) prior to its beginning, using user traffic statistics commonly provided by publishers and analytics companies. Nevertheless, our model can accommodate a random number of impressions through dummy arrivals that are valued at zero by the advertiser.

3 Off-equilibrium performance guarantees

In this section we study the performance that can be guaranteed by bidding strategies without any assumptions on the primitives (budgets, valuations) and strategies that govern competitors' bids. We obtain a lower bound on the minimal loss one must incur relative to the best dynamic response in hindsight under arbitrary competition. We then introduce a class of adaptive pacing strategies that react to realized valuations and expenditures throughout the campaign, and obtain an upper bound on the performance of these strategies under arbitrary and stationary competitive environments, thereby establishing asymptotic optimality in both settings.

3.1 Lower bound on the achievable guaranteed performance

Before investigating the type of performance an advertiser can aspire to, we first formalize the type of performance that *cannot* be achieved. The following result bounds the minimal loss one must incur relative to the best dynamic (and feasible) sequence of bids that could have been selected in hindsight, under arbitrary competition.

Theorem 3.1. (Lower bound on the achievable guaranteed performance) *For any target expenditure rate $\rho_k \in (0, \bar{v}_k]$ and number $\gamma < \bar{v}_k/\rho_k$ there exists some constant $\underline{C} > 0$ such that for any bidding strategy $\beta \in \mathcal{B}_k$:*

$$\limsup_{\substack{T \rightarrow \infty, \\ B_k = \rho_k T}} \sup_{\substack{\mathbf{v}_k \in [0, \bar{v}_k]^T \\ \mathbf{d}_k \in \mathbb{R}_+^T}} \frac{1}{T} \left(\pi_k^H(\mathbf{v}_k; \mathbf{d}_k) - \gamma \pi_k^\beta(\mathbf{v}_k; \mathbf{d}_k) \right) \geq \underline{C}.$$

Theorem 3.1 establishes that no admissible strategy can guarantee asymptotic γ -competitiveness for any $\gamma < \bar{v}_k/\rho_k$, and therefore cannot guarantee a portion larger than ρ_k/\bar{v}_k of the total net utility achieved by a dynamic response that is taken with the benefit of hindsight. Because an advertiser should never pay more than its value, \bar{v}_k is the maximal possible expenditure and, thus, the ratio ρ_k/\bar{v}_k captures the relative “wealthiness” of the advertiser. This implies that bidding strategies may perform poorly relative to best performance in hindsight when budgets are small,

but better performance might be guaranteed when budgets are bigger (when $\rho_k \geq \bar{v}_k$ the advertiser can bid truthfully to guarantee optimal performance and 1-competitiveness). Thus, Theorem 3.1 establishes the impossibility of "no-regret" relative to dynamic sequences of bids, by showing that, in general, the performance of any learning strategy cannot be guaranteed to approach π_k^H .

We next describe the main ideas in the proof of the theorem. We first observe, by adapting Yao's principle (Yao 1977) to our setting, that in order to bound the worst-case loss of any (deterministic or not) strategy relative to the best response in hindsight, it suffices to analyze the expected loss of deterministic strategies relative to that benchmark, where the sequence of competing bids is drawn from a certain distribution (see Lemma A.1). We next illustrate the idea behind the worst-case instance we construct. We assume that the advertiser knows in advance that the sequence of highest competing bids will be randomly selected from the set $\{\mathbf{d}^1, \mathbf{d}^2\}$, where

$$\mathbf{d}^1 = \left(\underbrace{d_{high}, \dots, d_{high}}_{\tau \text{ auctions}}, \underbrace{\bar{v}_k, \dots, \bar{v}_k}_{T-\tau \text{ auctions}} \right)$$

$$\mathbf{d}^2 = \left(\underbrace{d_{high}, \dots, d_{high}}_{\tau \text{ auctions}}, \underbrace{d_{low}, \dots, d_{low}}_{T-\tau \text{ auctions}} \right)$$

for some $\bar{v}_k \geq d_{high} > d_{low} > 0$, and that values will be \bar{v}_k throughout the campaign. We establish that in this case one may restrict analysis to strategies that determine prior to the beginning of the campaign how many auctions to win at different stages of the campaign (see Lemma A.2). This presents the advertiser with the following tradeoff: while early auctions introduce a return per unit of budget that is certain but low, later auctions introduce a return per unit of budget that may be higher but may also decrease to zero. Since the number of auctions in the first stage, denoted by τ , may be designed to grow with T , this tradeoff may drive a loss that does not diminish to zero asymptotically. The proof of the theorem presents a more general construction that follows the ideas illustrated above, and tunes the structural parameters to maximize the worst-case loss that must be incurred by any strategy relative to the best dynamic response in hindsight. Notably, as the above construction assumes that the advertiser's values are known and fixed to \bar{v}_k , Theorem 3.1 implies that, in general, no strategy can achieve a competitive ratio lower than \bar{v}_k/ρ_k (and, therefore, diminishing regret) even if advertisers know their own valuations.

3.2 Asymptotic optimal bidding strategy

We introduce an adaptive pacing strategy that adjust bids based on observations. In what follows we denote by $P_{[a,b]}(x) = \min\{\max\{x, a\}, b\}$ the Euclidean projection operator on the interval $[a, b]$.

Adaptive pacing strategy (A). Input: a number $\epsilon_k > 0$.

1. Select an initial multiplier $\mu_{k,1}$ in $[0, \bar{\mu}_k]$, and set the remaining budget to $\tilde{B}_{k,1} = B_k = \rho_k T$.
2. For each $t = 1, \dots, T$:
 - (a) Observe the realization of the random valuation $v_{k,t}$, and post a bid:

$$b_{k,t}^A = \min \left\{ v_{k,t}/(1 + \mu_{k,t}), \tilde{B}_{k,t} \right\}.$$

- (b) Observe the expenditure $z_{k,t}$. Update the multiplier by

$$\mu_{k,t+1} = P_{[0, \bar{\mu}_k]}(\mu_{k,t} - \epsilon_k (\rho_k - z_{k,t}))$$

and the remaining budget by $\tilde{B}_{k,t+1} = \tilde{B}_{k,t} - z_{k,t}$.

The adaptive pacing strategy dynamically adjusts the pace at which the advertiser depletes its budget by updating a multiplier that determines the extent to which the advertiser bids below its true values (shades bids). The strategy consists of a sequential approximation scheme that takes place in the Lagrangian dual space, and that is designed to approximate the best solution in hindsight defined in (1). The objective of the Lagrangian dual of (1) is $\sum_{t=1}^T x_{k,t}(v_{k,t} - (1 + \mu)d_{k,t}) + \mu\rho_k$, where μ is a nonnegative multiplier capturing the shadow price associated with the budget constraint. For a fixed μ , the dual objective is maximized by winning all items with $v_{k,t} \geq (1 + \mu)d_{k,t}$, which is obtained by bidding $b_{k,t} = v_{k,t}/(1 + \mu)$, since advertiser k wins whenever $b_{k,t} \geq d_{k,t}$. Therefore, one has

$$\pi_k^H(\mathbf{v}_k; \mathbf{d}_k) \leq \inf_{\mu \geq 0} \sum_{t=1}^T \underbrace{(v_{k,t} - (1 + \mu)d_{k,t})^+}_{=:\psi_{k,t}(\mu)} + \mu\rho_k, \quad (4)$$

where we denote by $y^+ = \max(y, 0)$ the positive part of a number $y \in \mathbb{R}$. The tightest upper bound can be obtained by solving the minimization problem on the right-hand side of (4). Since this problem cannot be solved without prior information on all the values and competing bids throughout the campaign, the proposed bidding strategy approximates (4) by estimating a multiplier $\mu_{k,t}$ and bidding $b_{k,t}^A = v_{k,t}/(1 + \mu_{k,t})$ in each round, as long as the remaining budget suffices.⁵

⁵The heuristic of bidding $b_{k,t} = v_{k,t}/(1 + \mu^H)$ until the budget is depleted, with μ^H optimal for (4), can be shown to be asymptotically optimal for the hindsight problem (1) as T grows large (see, e.g., Talluri and van Ryzin 1998).

The dual approximation scheme consists of estimating a direction of improvement and following that direction. More precisely, the sequence of multipliers follows a subgradient descent scheme, using the noisy point access each advertiser has to $\partial_- \psi_{k,t}(\mu_{k,t})$, namely, the left derivative of the t -period component of the dual objective in (4). In other words, in each period $t = 1, 2 \dots$, one has

$$\partial_- \psi_{k,t}(\mu_{k,t}) = \rho_k - d_{k,t} \mathbf{1}\{v_{k,t} \geq (1 + \mu_{k,t})d_{k,t}\} = \rho_k - d_{k,t} \mathbf{1}\{b_{k,t}^A \geq d_{k,t}\} = \rho_k - z_{k,t}.$$

In each period t the advertiser compares its expenditure $z_{k,t}$ to the target expenditure rate ρ_k that is “affordable” given the initial budget. Whenever the advertiser’s expenditure exceeds the target expenditure rate, $\mu_{k,t}$ is increased by $\epsilon_k(z_{k,t} - \rho_k)$, implying that the shadow price increases and that the associated budget constraint becomes more binding. On the other hand, if the advertiser’s expenditure is lower than the target expenditure rate (including periods in which the expenditure is zero) the multiplier $\mu_{k,t}$ is decreased by $\epsilon_k(\rho_k - z_{k,t})$, implying that the shadow price decreases and that the associated budget constraint becomes less binding. Then, given the valuation in that period, the advertiser shades its value using the current estimated multiplier.

The multiplier $\mu_{k,t}$ essentially captures a “belief” advertiser k has at time t regarding the shadow price associated with its budget constraint. Notably, the “correct” shadow price, reflecting the “correct” value of future opportunities, depends on unknown characteristics such as the value distribution of the advertiser, as well as the value distributions, budgets, and strategies of its competitors. Rather than aiming at estimating these unknown factors, the proposed strategy aims at learning the best response to these in terms of this shadow price.

Assumption 3.2. (Appropriate selection of step size) *The number ϵ_k satisfies $0 < \epsilon_k \leq 1/\bar{v}_k$,*

$$\lim_{T \rightarrow \infty} \epsilon_k = 0, \text{ and } \lim_{T \rightarrow \infty} T\epsilon_k = \infty.$$

Assumption 3.2 adapts a standard step size condition that is common in the stochastic approximation literature (see, e.g., §1.1 in Kushner and Yin 2003) to stationary step sizes, and characterizes the range of rates at which step sizes should converge to zero (as a function of the horizon length) to allow for the convergence of the scheme, while guaranteeing sufficient exploration. By doing so, this assumption prescribes a broad range for tuning the strategy’s step size in a manner that guarantees the asymptotic optimality results that follow.⁶ Even though Assumption 3.2 places

⁶We restrict the formal definition of the adaptive pacing strategy to stationary step sizes only to simplify and shorten analysis, and the strategy can be adjusted to allow for time-varying step sizes through a more general form of Assumption 3.2. Nevertheless, we note that the asymptotic optimality notions we consider are defined for the general class of admissible budget-feasible strategies and achieved within the subclass of strategies with stationary step sizes.

some restrictions on the step size selection, it still defines a broad class of step sizes; for example, $\epsilon_k = c/T^\gamma$ satisfies Assumption 3.2 for any constants $0 < c \leq 1/\bar{v}_k$ and $0 < \gamma < 1$.

3.3 Performance under arbitrary competition

The following result characterizes the performance that can be guaranteed by an adaptive pacing strategy under arbitrary competition.

Theorem 3.3. (Asymptotic optimality under arbitrary competition) *Let A be an adaptive pacing strategy with a step size ϵ_k satisfying Assumption 3.2 and $\bar{\mu}_k \geq \bar{v}_k/\rho_k - 1$. Then:*

$$\limsup_{\substack{T \rightarrow \infty, \\ B_k = \rho_k T}} \sup_{\substack{\mathbf{v}_k \in [0, \bar{v}_k]^T \\ \mathbf{d}_k \in \mathbb{R}_+^T}} \frac{1}{T} \left(\pi_k^H(\mathbf{v}_k, \mathbf{d}_k) - \frac{\bar{v}_k}{\rho_k} \pi_k^A(\mathbf{v}_k, \mathbf{d}_k) \right) = 0.$$

Moreover, when the step size ϵ_k is of order $T^{-1/2}$, the convergence is at rate $T^{-1/2}$.

Theorem 3.3 establishes that an adaptive pacing strategy is asymptotic (\bar{v}_k/ρ_k) -competitive (Borodin and El-Yaniv 1998), guaranteeing at least a fraction ρ_k/\bar{v}_k of the total net utility achieved by any dynamic response that could have been selected with the benefit of hindsight. This establishes the asymptotic optimality of the strategy, as by Theorem 3.1 no other admissible bidding strategy can guarantee a larger fraction of the best performance in hindsight. (Recalling the worst-case instance used in Theorem 3.1, a larger fraction of the best performance in hindsight cannot be guaranteed even when advertisers know their valuations up front.)

In the proof of the theorem, we first show that the time at which the budget is depleted under the adaptive pacing strategy is “close” to T . We then express the loss incurred by that strategy relative to the best response in hindsight in terms of the value of lost auctions, and bound the potential value of these auctions. Together, these analyses establish that there exist some positive constants C_1, C_2 , and C_3 , independent of T , such that for any $\mathbf{d}_k \in \mathbb{R}_+^T$ and $\mathbf{v}_k \in [0, \bar{v}_k]^T$,

$$\pi_k^H(\mathbf{v}_k; \mathbf{d}_k) - \frac{\bar{v}_k}{\rho_k} \pi_k^A(\mathbf{v}_k, \mathbf{d}_k) \leq C_1 + \frac{C_2}{\epsilon_k} + C_3 T \epsilon_k.$$

This allows one to evaluate the performance of the various step sizes, and establishes asymptotical optimality whenever Assumption 3.2 holds (independently of the initial multiplier $\mu_{1,k}$). In particular, selecting $\epsilon_k = cT^{-1/2}$ with $c = (C_2/C_3)^{1/2}$ leads to a convergence rate of $T^{-1/2}$.

3.4 Performance under stationary competition

The upper bound in Theorem 3.3 holds in every sample path (as the adaptive pacing strategy is deterministic) and is extremely robust. For example, it holds also when campaigns can start and finish at different times and under arbitrary sequences of valuations and competing bids, including sequences that are correlated across impressions or advertisers, as well as dynamic sequences that are based on knowledge about the bidding strategy and its realized path. However, the established performance guarantee could be viewed as conservative, or “pessimistic,” in the sense that it holds under *any* sequences of valuations and competing bids, including worst-case instances such as the one described in the proof of Theorem 3.1. In particular, when the target expenditure rate ρ_k is small relative to \bar{v}_k , the guaranteed performance is limited as well. In this subsection we demonstrate that for any fixed target expenditure rate, the asymptotic optimality of an adaptive pacing strategy can also be carried over (with much better performance) to more “optimistic” settings in which uncertainties are drawn independently from a stationary distribution. (See Iyer et al. 2014; Balseiro et al. 2015 for discussions on the validity of assuming stationarity competition.)

We assume that the pairs $(v_{k,t}, d_{k,t})$ are independently drawn from some unknown stationary distribution. We denote by $\Psi_k(\mu) := \sum_{t=1}^T \left(\mathbb{E}_{v_{k,t}, d_{k,t}} [(v_{k,t} - (1 + \mu)d_{k,t})^+] + \mu\rho_k \right)$ the expectation of the dual objective function given in (4), with respect to values and competing bids.

Theorem 3.4. (Asymptotic optimality under stationary competition) *Suppose that in each period t the pair $(v_{k,t}, d_{k,t})$ is independently drawn from some stationary distribution such that the dual function $\Psi_k(\mu)$ is thrice differentiable with bounded derivatives and strongly convex with parameter $\lambda_k > 0$. Let A be an adaptive pacing strategy with a step size ϵ_k satisfying Assumption 3.2, as well as $\epsilon_k < 1/(2\lambda_k)$ and $\bar{\mu}_k \geq \bar{v}_k/\rho_k - 1$. Then:*

$$\limsup_{\substack{T \rightarrow \infty, \\ B_k = \rho_k T}} \frac{1}{T} \mathbb{E}_{\mathbf{v}_k, \mathbf{d}_k} [\pi_k^H(\mathbf{v}_k, \mathbf{d}_k) - \pi_k^A(\mathbf{v}_k, \mathbf{d}_k)] = 0.$$

Moreover, when the step size ϵ_k is of order $T^{-1/2}$, the convergence is at rate $T^{-1/2}$.

Theorem 3.4 establishes that when valuations and competing bids are independent and identically drawn throughout time, an adaptive pacing strategy is long-run average optimal, in the sense that for *any* fixed $\rho_k > 0$, its average expected performance converges to the average expected performance of the best sequence of bids that could have been selected in hindsight. Together with Theorems 3.1 and 3.3, Theorem 3.4 establishes the asymptotic optimality of the adaptive

pricing strategy in both the “pessimistic” and “optimistic” scenarios when advertisers have *no prior knowledge* of whether the nature of the competition is “stationary” or “arbitrary,” and that selecting $\epsilon_k \sim T^{-1/2}$ guarantees a convergence rate of $T^{-1/2}$ in both frameworks. We leave the characterization of rate optimality in these frameworks as an open problem.

Theorem 3.4 is based on the assumption that the dual function is thrice differentiable and strongly convex. The first condition is technical and is required to perform Taylor series expansions, and the second condition is a standard one guaranteeing that the multipliers employed by an adaptive pricing strategy converge. Lemma C.1 in Appendix C shows that these conditions hold when values and competing bids are independent and absolutely continuous with bounded densities and the valuation density is differentiable. The conditions in Theorem 3.4 can also be shown to hold when competing bids are (imperfectly) correlated across auctions.

To prove the theorem we first upper bound the expected performance in hindsight in terms of $\Psi(\mu^*)$, where μ^* minimizes the dual function. We then lower bound the performance of the adaptive pricing strategy in terms of $\Psi(\mu^*)$ by developing a second-order expansion around μ^* of the expected utility per auction when the strategy employs a multiplier $\mu_{k,t}$. This expansion involves the mean error $\mathbb{E}[\mu_{k,t} - \mu^*]$ in the first-order term and the mean squared error $\mathbb{E}[(\mu_{k,t} - \mu^*)^2]$ in the second-order term. Using the update rule we control the cumulative sum of these errors to establish that there exist some positive constants C_1 , C_2 , and C_3 , independent of T , such that

$$\mathbb{E}_{\mathbf{v}_k, \mathbf{d}_k} [\pi_k^H(\mathbf{v}_k, \mathbf{d}_k) - \pi_k^A(\mathbf{v}_k, \mathbf{d}_k)] \leq C_1 + \frac{C_2}{\epsilon_k} + C_3 T \epsilon_k.$$

This establishes asymptotic optimality whenever the sequence of step sizes satisfies Assumption 3.2 and establishes a convergence rate of $T^{-1/2}$ under a step size selection of order $T^{-1/2}$.

3.5 Empirical proof of concept using ad auction data

To analyze the value that may be captured in practice by an adaptive pricing strategy, we empirically evaluate the portion this strategy recovers of the best performance in hindsight, using data from a large online ad auction platform. The data set includes bids from 13,182 sequential online ad auctions, in which a total of 2,923 advertisers participated. The objective of the following analysis is to demonstrate the effectiveness of an adaptive pricing strategy in a realistic yet simple setting. In particular, an optimal tuning of the strategy’s parameters may depend on the available budget as well as competition characteristics. Obtaining the optimal characterization of these parameters is a challenging open problem that we leave as future work.

Setup. For each advertiser k , we set the sequence of valuations \mathbf{v}_k to be the sequence of recorded bids submitted by that advertiser (where \bar{v}_k is the highest recorded bid of advertiser k), and the sequence \mathbf{d}_k to be the sequence of highest bids submitted by other advertisers. We randomly generated $\rho_k \in [0, \bar{v}_k]$ according to a uniform distribution, and for each target expenditure rate we approximated $\pi_k^H(\mathbf{v}_k, \mathbf{d}_k)$ using a standard linear programming relaxation. Note that whenever an advertiser is effectively unconstrained, that is, $\sum_{t=1}^T \mathbf{1}\{d_{k,t} \leq v_{k,t}\} d_{k,t} \leq \rho_k T$, the best performance in hindsight, π_k^H , can be implemented by truthful bidding. We then computed the performance of an adaptive pacing strategy π_k^A , and compared its performance with π_k^H . We focused on the 10 bidders with the largest cumulative expenditure (together, these advertisers were responsible for more than 98% of the overall expenditure), and for each of these advertisers we replicated the process 200 times. In all the instances we used an initial multiplier $\mu_1 = 0.5$, and a step size $\epsilon = T^{-1/2}$ (the choice of the initial multiplier had very limited impact on the performance).

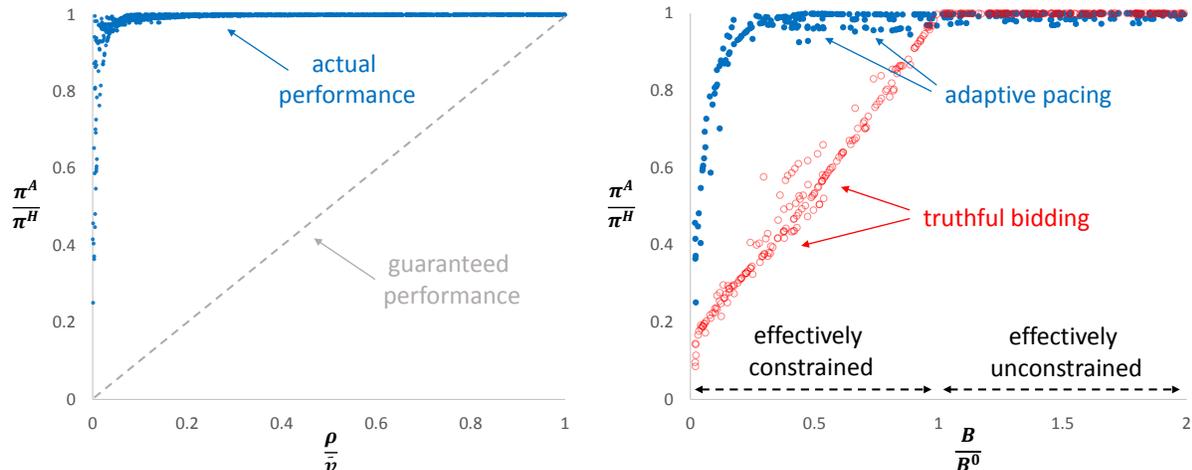


Figure 1: Performance of an adaptive pacing strategy with parameters $\mu_1 = 0.5$ and $\epsilon = T^{-1/2}$, in terms of the attained fractions of the performance of the best sequence of bids in hindsight. (*Left*) Compared with the fractions that are guaranteed under arbitrary competition, as a function of ρ/\bar{v} . (*Right*) Compared with truthful bidding, as a function of B/B^0 .

Results and discussion. The plot that appears at the left-hand side of Figure 1 depicts the fractions of π_k^H (as a function of ρ_k/\bar{v}_k) that were recovered by the adaptive pacing strategy. Notably, in all the instances the strategy significantly outperformed the payoff guarantees for an arbitrary competitive environment. Moreover, it approximately achieved π_k^H , except when the target expenditure rate was particularly small. The plot on the right-hand side of Figure 1 compares the fractions of π_k^H that were recovered by the adaptive pacing strategy with those recovered by “truthful bidding,” a naïve strategy that bids the advertiser’s value until the budget is depleted (each dot corresponds to

the performance of one advertiser under one target expenditure rate). We plotted these quantities as a function of B_k/B_k^0 , where $B_k = \rho_k T$ is the available budget and $B_k^0 = \sum_{t=1}^T \mathbf{1}\{d_{k,t} \leq v_{k,t}\} d_{k,t}$ is the minimal budget under which the advertiser is effectively unconstrained and truthful bidding is optimal (the plot focuses on instances with $B_k/B_k^0 \leq 2$). One may expect that whenever $\rho_k < \bar{v}_k$ but $B_k/B_k^0 \geq 1$, any non-anticipating strategy would incur a loss (relative to truthful bidding) that is associated with not having ex ante knowledge of being effectively unconstrained. However, in these cases the loss incurred by the adaptive pacing strategy is less than 1%. On the other hand, whenever $B_k/B_k^0 < 1$, the advertiser is effectively constrained. On average, in these cases the adaptive pacing strategy outperforms truthful bidding by 120%. We attribute the relative deterioration in the empirical performance of adaptive pacing for small values of ρ to the fact that, for fixed valuations and competing bids, a very small budget suffices for winning only a few auctions (under such conditions one could expect the ability of learning policies to compete with the best sequence of bids in hindsight to be limited). We further note that one can likely improve the empirical performance of the strategy by customizing the selection of the tuning parameters for different advertisers. Overall, these results demonstrate the practical effectiveness of an adaptive pacing strategy even under moderate budgets and problem horizons, as well as realistic bidding behavior that may be subject to correlated preferences.

4 Convergence of simultaneous learning

We turn to study the dynamics that emerge when all advertisers follow adaptive pacing strategies simultaneously. We establish that the sequences of multipliers generated by the strategies converge to a tractable profile of multipliers, and that the payoffs converge to the payoffs achieved under such a profile.

4.1 A candidate profile of dual multipliers

Under simultaneous adoption of adaptive pacing strategies, each advertiser follows the sub-gradient $\partial_- \psi_{k,t}(\mu_{k,t}) = \rho_k - z_{k,t}$. Therefore, if the induced sequences of multipliers converge, one may expect the limiting profile to consist of multipliers under which each advertiser's expected expenditure equals its target expenditure (whenever its budget is binding). Denote the expected expenditure per auction of advertiser k when advertisers shade bids according to a profile of multipliers $\boldsymbol{\mu} \in \mathbb{R}^K$ by $G_k(\boldsymbol{\mu}) := \mathbb{E}_{\mathbf{v}} [\mathbf{1}\{(1 + \mu_k)d_k \leq v_k\} d_k]$, where $d_k = \max_{i:i \neq k} \left\{ v_i / (1 + \mu_i) \right\}$ and the expectation is taken w.r.t. the values $v_k \sim F_k$. Following the above intuition, we consider the vector $\boldsymbol{\mu}^*$, defined

by the complementary conditions:

$$\mu_k^* \geq 0 \quad \perp \quad G_k(\boldsymbol{\mu}^*) \leq \rho_k, \quad \forall k = 1, \dots, K, \quad (5)$$

where \perp indicates a complementarity condition between the multiplier and the expenditure; that is, at least one condition must hold with equality. Intuitively, each advertiser shades bids if its expected expenditure is equal to its target rate. Conversely, if an advertiser's target rate exceeds its expected expenditure, then it bids truthfully by setting the multiplier equal to zero.

We next provide sufficient conditions for the uniqueness of the vector $\boldsymbol{\mu}^*$. The vector function $\mathbf{G} : \mathbb{R}^K \rightarrow \mathbb{R}^K$ is said to be λ -strongly monotone over a set $\mathcal{U} \subset \mathbb{R}^K$ if $(\boldsymbol{\mu} - \boldsymbol{\mu}')^\top (\mathbf{G}(\boldsymbol{\mu}') - \mathbf{G}(\boldsymbol{\mu})) \geq \lambda \|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_2^2$ for all $\boldsymbol{\mu}, \boldsymbol{\mu}' \in \mathcal{U}$. We refer to λ as the monotonicity constant.

Assumption 4.1. (Stability)

1. *There exists a set $\mathcal{U} = \prod_{k=1}^K [0, \bar{\mu}_k]$ and a monotonicity constant $\lambda > 0$, independent of K , such that the expected expenditure function $\mathbf{G}(\boldsymbol{\mu})$ is λ -strongly monotone over \mathcal{U} .*
2. *The target expenditure rate satisfies $\rho_k \geq \bar{v}_k / \bar{\mu}_k$ for every bidder k .*

The first part of Assumption 4.1 requires the expenditure function to be strongly monotone over a set of feasible multipliers. In the case of a single agent, the strong monotonicity of the expenditure function is equivalent to the strong convexity of the dual objective function (see Theorem 3.4). Therefore, the monotonicity condition is a natural extension of the strong convexity condition to a simultaneous learning setting. The monotonicity condition is similar to other common conditions that guarantee uniqueness and stability in concave (unconstrained) games; see, e.g., Rosen (1965), as well as the discussion on convergence of stochastic approximation methods in multi-dimensional action spaces in §1.10, Part 2 of Benveniste et al. (1990). To provide some intuition on this condition, consider a scenario in which the expected expenditure per auction of all advertisers is above their target expenditure rate (that is, $G_k(\boldsymbol{\mu}) > \rho_k$ for all k). Because the expenditure function G_k is non-increasing in μ_k for all k , each advertiser can get closer to its target expenditure rate by unilaterally increasing the multiplier by a small amount. This condition guarantees that if *all advertisers* simultaneously increase their multipliers by a small amount, then all advertisers get closer to their target expenditure rates.⁷ The second part of the assumption is that budgets are not too small and guarantees that $\boldsymbol{\mu}^* \in \mathcal{U}$. This part suggests a simple rule of thumb for

⁷In Appendix C we show that this part of the assumption is implied by the diagonal strict concavity condition defined in Rosen (1965), prove that it holds in symmetric settings, and demonstrate its validity in simple cases.

setting a sufficiently large upper bound $\bar{\mu}_k$. In Appendix C (see Proposition C.3) we show that Assumption 4.1 guarantees the existence of a unique vector of multipliers $\boldsymbol{\mu}^*$. As we further discuss in §6, the uniqueness of $\boldsymbol{\mu}^*$ has practical benefits that may motivate a market designer to place restrictions on the target expenditure rates of advertisers to guarantee that Assumption 4.1 holds.

4.2 Convergence of dual multipliers

To analyze the sample path of multipliers resulting from simultaneous adoption of adaptive pacing strategies, we consider some conditions on the step sizes selected by the different advertisers.

Assumption 4.2. (Joint selection of step sizes) *Let $\bar{\epsilon} = \max_{k \in \{1, \dots, K\}} \epsilon_k$, $\underline{\epsilon} = \min_{k \in \{1, \dots, K\}} \epsilon_k$, $\bar{v} = \max_k \{\bar{v}_k\}$, and let λ be the monotonicity constant that appears in Assumption 4.1. The profile of step sizes $(\epsilon_k)_{k=1}^K$ satisfies:*

- (i) $\bar{\epsilon} \leq 1/\bar{v}$ and $\bar{\epsilon} < 1/(2\lambda)$,
- (ii) $\lim_{T \rightarrow \infty} \bar{\epsilon}^2/\underline{\epsilon} = 0$,
- (iii) $\lim_{T \rightarrow \infty} T\underline{\epsilon}^2/\bar{\epsilon} = \infty$.

Assumption 4.2 details conditions on the *joint* selections of step sizes (extending the individual conditions that appear in Assumption 3.2), essentially requiring the step sizes selected by the different advertisers to be “reasonably close” to one another. Nevertheless, the class of step sizes defined by Assumption 4.2 is quite general and allows for flexibility in the individual step size selection. For example, step sizes $\epsilon_k = c_k T^{-\gamma_k}$ satisfy Assumption 4.2 for any constants $c_k > 0$ and $0 < \gamma_k < 1$ such that $\max_k \gamma_k \leq 2 \min_k \gamma_k$ and $2 \max_k \gamma_k \leq 1 + \min_k \gamma_k$. Moreover, when step sizes are equal across advertisers and satisfy $\epsilon < 1/(2\lambda)$, Assumption 4.2 coincides with Assumption 3.2.

The following result states that, when Assumption 4.2 holds, the multipliers selected by advertisers under simultaneous adoption of adaptive pacing strategies converge to the vector $\boldsymbol{\mu}^*$. In what follows we use the standard norm notation $\|\mathbf{x}\|_p := \left(\sum_{k=1}^K |x_k|^p\right)^{1/p}$ for a vector $\mathbf{x} \in \mathbb{R}^K$.

Theorem 4.3. (Convergence of dual multipliers) *Suppose that Assumption 4.1 holds, and let $\boldsymbol{\mu}^*$ be the profile multipliers defined by (5). Assume that all advertisers follow adaptive pacing strategies with step sizes that together satisfy Assumption 4.2. Then:*

$$\lim_{\substack{T \rightarrow \infty, \\ B_k = \rho_k T}} \frac{1}{T} \sum_{t=1}^T \mathbb{E} [\|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2] = 0.$$

Theorem 4.3 adapts standard stochastic approximation techniques (see, e.g., Nemirovski et al. 2009) to accommodate the expenditure feedback form and the possibility of different step sizes across agents. In the proof of the theorem we establish that under Assumption 4.1 there exist positive constants C_1 and C_2 , independent of T , such that

$$\mathbb{E} \left[\|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2 \mathbf{1} \left\{ \tilde{B}_{k,t+1} \geq \bar{v}_k \quad \forall k \right\} \right] \leq C_1 \frac{\bar{\epsilon}}{\underline{\epsilon}} (1 - 2\lambda\underline{\epsilon})^{t-1} + C_2 \frac{\bar{\epsilon}^2}{\underline{\epsilon}},$$

for any $t \in \{1, \dots, T\}$, where the dependence of C_1 and C_2 on the number of advertisers is of a linear order. This implies that insofar as the respective primal solutions are feasible given the remaining budgets, the vector of multipliers $\boldsymbol{\mu}_t$ converges in L_2 to the vector $\boldsymbol{\mu}^*$. We then argue that the strategies do not deplete the budget too early, and establish the result using the conditions in Assumption 4.2. We note that the fastest rate of convergence within this class of strategies is of order $T^{-1/2}$, and is achieved under the symmetric step-size selection of the form $\epsilon \sim T^{-1/2}$.

Due to primal feasibility constraints, the convergence of multipliers does not guarantee convergence in performance. However, the former has stand-alone practical importance in allowing the market designer to predict the bidding behavior of advertisers. This implication is discussed in §6.

4.3 Convergence in performance

We next study the long-run average performance achieved when all advertisers follow adaptive pacing strategies. Recall that Theorem 4.3 implies that the vector of multipliers $\boldsymbol{\mu}_t$ selected by advertisers under these strategies converges to the vector $\boldsymbol{\mu}^*$. However, since in every period the remaining budget may potentially constrain bids, the convergence of multipliers to $\boldsymbol{\mu}^*$ does not guarantee performance that is close to the performance achieved when using the fixed vector $\boldsymbol{\mu}^*$ throughout the campaign. Nevertheless, establishing convergence in performance is a key step toward equilibrium analysis, and the performance that is obtained using the fixed vector $\boldsymbol{\mu}^*$ is a natural candidate for such a convergence. Given a vector of multipliers $\boldsymbol{\mu}$, we denote

$$\Psi_k(\boldsymbol{\mu}) := T \left(\mathbb{E}_{\mathbf{v}} \left[(v_k - (1 + \mu_k)d_k)^+ \right] + \mu_k \rho_k \right), \quad (6)$$

where d_k is defined as in §4.1. Recalling (4), $\Psi_k(\boldsymbol{\mu}^*)$ is the dual performance of advertiser k when in each period t each advertiser i bids $b_{i,t} = v_{i,t}/(1 + \mu_i^*)$. Balseiro et al. (2015) show that $\Psi_k(\boldsymbol{\mu}^*)$ coincides with the expected performance achieved when all advertisers shade bids according to $\boldsymbol{\mu}^*$ and are allowed to bid even after budgets are depleted.

Theorem 4.4. (Convergence in performance) *Suppose that Assumption 4.1 holds, and let $\boldsymbol{\mu}^*$ be the profile of multipliers defined by (5). Let $\mathbf{A} = (\mathbf{A}_k)_{k=1}^K$ be a profile of adaptive pacing strategies, with step sizes that together satisfy Assumption 4.2. Then, for each $k \in \{1, \dots, K\}$:*

$$\limsup_{\substack{T \rightarrow \infty, \\ B_k = \rho_k T}} \frac{1}{T} (\Psi_k(\boldsymbol{\mu}^*) - \Pi_k^{\mathbf{A}}) \leq 0.$$

In the proof we establish that when all advertisers follow adaptive pacing strategies, deviations from target expenditure rates are controlled in a manner such that advertisers very often deplete their budgets “close” to the end of the campaign. Therefore the potential loss along the sample path due budget depletion is relatively small compared to the cumulative payoff. We establish that there exist some positive constants C_1 through C_4 , independent of T , such that

$$\Psi_k(\boldsymbol{\mu}^*) - \Pi_k^{\mathbf{A}} \leq C_1 + C_2 \frac{\bar{\epsilon}^{1/2}}{\underline{\epsilon}^{3/2}} + C_3 \frac{\bar{\epsilon} T}{\underline{\epsilon}^{1/2}} + C_4 \frac{1}{\underline{\epsilon}}. \quad (7)$$

This allows for the analysis of various step size selections, and establishes convergence in performance whenever the sequence of step sizes satisfies Assumption 4.2. In particular, (7) implies that a symmetric step-size selection of order $\epsilon \sim T^{-1/2}$ guarantees a long-run average convergence rate of order $T^{-1/4}$. We conjecture that such parametric selection further guarantees a convergence rate of $T^{-1/2}$, and demonstrate this numerically Appendix D.

Illustrative numerical analysis of convergence under simultaneous learning. In Appendix D we demonstrate the convergence established in Theorems 4.3 and 4.4 through a sequence of numerical experiments that simulate the sample path and performance of bidders that simultaneously follow adaptive pacing strategies. Our results demonstrate the asymptotic convergence of dual multipliers and performance in a variety of instance and under different selections of step sizes that satisfy Assumption 4.2. A step size selection of $\epsilon = T^{-1/2}$ led to superior convergence rates of $T^{-1/2}$, for both multipliers and performance.

5 Approximate Nash equilibrium in dynamic strategies

In this section we study conditions under which adaptive pacing strategies constitute an ϵ -Nash equilibrium. First, we note that while a profile of adaptive pacing strategies converges in multipliers and in performance, such a strategy profile does not necessarily constitute an approximate Nash equilibrium. We then detail one extension of our formulation (motivated in the context of ad

auctions) where in each period each advertiser participates in one of multiple auctions that take place simultaneously. We show that in such a regime adaptive pacing strategies constitute an ε -Nash equilibrium in dynamic strategies.

An advertiser can potentially gain by deviating from a profile of adaptive pacing strategies if this deviation: (i) reduces competition in later auctions by causing competitors to deplete their budgets earlier; or (ii) reduces competition before the budgets of the competitors are depleted by inducing competitors to learn “wrong” multipliers. Under Assumption 4.2, the self-correcting nature of adaptive pacing strategies prevents competitors from depleting budgets too early and thus the first avenue is not effective. However, when the same set of advertisers frequently interact, an advertiser can cause competitors to excessively shade bids by driving them to learn higher multipliers and, therefore, can reduce competition and increase profits throughout the campaign.

In practice, however, the number of advertisers bidding in an online ad market is typically large and, due to the high frequency of auctions as well as sophisticated targeting technologies, advertisers often participate only in a fraction of all auctions that take place (Celis et al. 2014). As a result, each advertiser often competes with different bidders in every auction (Balseiro et al. 2014). This limits the impact a single advertiser can have on the expenditure trajectory of other advertisers and, therefore, the benefit from unilaterally deviating to strategies such as the ones described above can be expected to be small. To demonstrate this intuition, we next detail one extension of our basic framework, where in each period multiple auctions take place simultaneously and each advertiser is “matched” with one of these auctions based on its targeting criteria and the attributes associated with the auctioned impressions. We establish that in markets with such characteristics adaptive pacing strategies constitute an ε -Nash equilibrium.

Parallel auctions model. In each time period $t = 1, \dots, T$ we now assume that there are M ad slots sold simultaneously in parallel auctions. This can correspond to multiple publishers simultaneously auctioning inventory or a publisher with multiple ad slots running one auction per slot. We assume that each advertiser participates in one of the M auctions independently at random. We denote by $m_{k,t} \in \{1, \dots, M\}$ the auction advertiser k participates in at time t . In each period t the random variable $m_{k,t}$ is independently drawn from an exogenous distribution $\alpha_k = \{\alpha_{k,m}\}_{m=1}^M$, where $\alpha_{k,m}$ denotes the probability that advertiser k participates in auction m . The type of advertiser k is therefore given by $\theta_k := (F_k, \rho_k, \alpha_k)$. The definition of $d_{k,t}$, the highest competing bid faced by advertiser k , is adjusted to be

$$d_{k,t} := \max \left\{ \max_{i: i \neq k} \left\{ \mathbf{1}\{m_{i,t} = m_{k,t}\} b_{i,t} \right\}, 0 \right\}.$$

Notably, all the results we established thus far hold when advertisers compete in M parallel auctions as discussed above. We next establish that under this extension adaptive pacing strategies constitute an ε -Nash equilibrium in dynamic strategies, in the sense that no advertiser can benefit from unilaterally deviating to any other strategy (including ones with access to complete information) when the number of time periods and bidders is large. We note that by applying a similar analysis, these strategies could also be shown to constitute an ε -Nash equilibrium in dynamic strategies under other frameworks that have been considered in the literature, such as those in Iyer et al. (2014) and Balseiro et al. (2014), in which advertisers' campaigns do not necessarily start and end at the same time. Additionally, the results of this section generalize to more sophisticated matching models, such as ones that allow the advertisers to participate in multiple slots per period (when the number of slots each advertiser participates in per period is small relative to M), and the distribution of values to be auction dependent.

We denote by $\mathcal{B}_k^{\text{CI}} \supseteq \mathcal{B}_k$ the space of non-anticipating and budget-feasible strategies with complete information. In particular, strategies in $\mathcal{B}_k^{\text{CI}}$ may have access to all types $(\theta_i)_{i=1}^K$, as well as past value realizations, bids, and expenditures of all competitors. To analyze unilateral deviations from a profile of adaptive pacing strategies we consider conditions on the likelihood that different advertisers compete in the same auction in a given time period. Let $a_{k,i} = \mathbb{P}\{m_{k,t} = m_{i,t}\} = \sum_{m=1}^M \alpha_{k,m} \alpha_{i,m}$ denote the probability that advertisers k and i compete in the same auction in a given time period, and let $\mathbf{a}_k = (a_{k,i})_{i \neq k} \in [0, 1]^{K-1}$. We refer to these as the *matching probabilities* of advertiser k .

Assumption 5.1. (Limited interaction) *The matching probabilities $(\mathbf{a}_k)_{k=1}^K$ and profile of step sizes $(\epsilon_k)_{k=1}^K$ satisfy:*

- (i) $\lim_{M, K \rightarrow \infty} K^{1/2} \max_{k=1, \dots, K} \|\mathbf{a}_k\|_2 < \infty,$
- (ii) $\lim_{M, K, T \rightarrow \infty} \bar{\epsilon}/\underline{\epsilon} \max_{k=1, \dots, K} \|\mathbf{a}_k\|_2^2 = 0.$

The first condition on the matching probabilities in Assumption 5.1 is that each advertiser interacts with a limited number of different competitors per time period, even when the total number of advertisers and auctions per period is large. The second condition ties the matching probabilities with the profile of step sizes selected by advertisers, essentially implying that advertisers can interact with more competitors per time period when the “learning rates” of the strategies are similar to one another. The class of matching probabilities and step sizes defined by Assumption 5.1 includes

many practical settings. For example, when each advertiser participates in each auction with the same probability, one has $\|\mathbf{a}_k\|_2 \approx K^{1/2}/M$ since $\alpha_{i,m} = 1/M$ for each advertiser i and auction m . When the expected number of bidders per auction $\kappa := K/M$ is fixed (implying that the number of parallel auctions is proportional to the number of players), we have $\|\mathbf{a}_k\|_2 \approx \kappa/K^{1/2}$ and the first condition is satisfied. The second condition states that in such a case the difference between the learning rates of different advertisers should be of order $o(K/\kappa^2)$.

Theorem 5.2. (ε -Nash equilibrium in dynamic strategies) *Suppose that Assumptions 4.1 and 5.1 hold. Let \mathbf{A} be a profile of adaptive pacing strategies with step sizes that together satisfy Assumption 4.2. Then:*

$$\limsup_{\substack{T, K, M \rightarrow \infty, \\ B_k = \rho_k T}} \sup_{\substack{k \in \{1, \dots, K\}, \\ \beta \in \mathcal{B}_k^{\text{CI}}}} \frac{1}{T} \left(\Pi_k^{\beta, \mathbf{A}-k} - \Pi_k^{\mathbf{A}} \right) \leq 0.$$

Theorem 5.2 establishes that adaptive pacing strategies (with step sizes and matching probabilities that together satisfy Assumptions 4.2 and 5.1) constitute an ε -Nash equilibrium within the class of dynamic strategies with access to complete information on market primitives, and perfect information on all the events that took place in prior periods. In particular, the benefit from unilaterally deviating from a profile of adaptive pacing strategies diminishes to zero in large markets, even when advertisers know their value distribution up front (as well as the value distributions and budgets of their competitors), and even when they can acquire real-time information on the past value realizations, bids, and payoffs of their competitors.

The key idea of the proof lies in bounding the benefit in terms of $\Psi_k(\boldsymbol{\mu}^*)$ from unilaterally deviating to an arbitrary strategy. Under Assumption 5.1 the impact of advertisers on one another is limited since each advertiser interacts with a limited number of different competitors in each time period. Thus, competitors learn the stable multipliers $\boldsymbol{\mu}^*$ regardless of the actions of a deviating advertiser and an advertiser cannot reduce competition by inducing competitors to learn a “wrong” multiplier. We show that the benefit from deviating to any other strategy is small when multipliers are “close” to $\boldsymbol{\mu}^*$. We bound the performance of an arbitrary budget-feasible strategy using a Lagrangian relaxation in which we add the budget constraint to the objective with μ_k^* as the Lagrange multiplier. This yields a bound on the performance of any dynamic strategy, including strategies with access to complete and perfect information. Thus, we obtain that there exist constants C_1 through C_5 , independent of T , K , and M , such that for any profile of advertisers satisfying Assumption 4.1 and any strategy $\beta \in \mathcal{B}_k^{\text{CI}}$:

$$\Pi_k^{\beta, \mathbf{A}_k} - \Psi_k(\boldsymbol{\mu}^*) \leq C_1 + C_2 \|\mathbf{a}_k\|_2 K^{1/2} \frac{\bar{\epsilon}^{1/2}}{\underline{\epsilon}^{3/2}} + C_3 \|\mathbf{a}_k\|_2 K^{1/2} \frac{\bar{\epsilon} T}{\underline{\epsilon}^{1/2}} + C_4 \frac{1}{\underline{\epsilon}} + C_5 \|\mathbf{a}_k\|_2^2 \frac{\bar{\epsilon} T}{\underline{\epsilon}}$$

for each advertiser k . This allows one to evaluate unilateral deviations under different step sizes, and to establish the result under the conditions in Assumptions 4.2 and 5.1.

Theorem 5.2 shows that a profile of adaptive pacing strategies constitutes an approximate equilibrium when there is a large number of players who do not interact too frequently. In some markets, however, a small group of advertisers may interact with each other in many auctions, thereby violating the assumptions of the theorem. Using a fluid model it is possible to numerically analyze the benefit from unilaterally deviating from a profile of adaptive pacing strategies in markets with a few advertisers. In numerical experiments we observed that the suboptimality of adaptive pacing strategies diminishes fast as the number of players increases. In particular, the benefit from deviating is typically small even in markets with few advertisers who interact frequently.

6 Concluding remarks

Summary and implications. In this paper we introduced a family of adaptive pacing strategies, in which advertisers adjust the pace at which they spend their budget according to their expenditures. We established that such strategies are asymptotically optimal when the competitive environment is stationary, but also when competitors' bids are arbitrary. We further demonstrated that in large markets, these strategies constitute an approximate Nash equilibrium in dynamic strategies. Notably, adaptive pacing strategies maintain these notions of optimality without requiring any prior knowledge of whether the nature of competition is stationary, arbitrary, or simultaneously learning. This implies that given characteristics that are common in online ad markets, an advertiser can essentially follow an approximate equilibrium bidding strategy while ensuring the best performance that can be guaranteed off equilibrium.

While we focus is on the advertisers' perspective, our results have implications on market design as well. Given access to the advertisers' types, our convergence results provide practical tools for predicting the bidding behavior of advertisers (through the vector $\boldsymbol{\mu}^*$) and their payoff (through the function Ψ). Since these predictions rely on the uniqueness of $\boldsymbol{\mu}^*$, this motivates the market designer to place restrictions on target expenditure rates that satisfy Assumption 4.1. In addition, our equilibrium analysis indicates that platforms (such as Facebook, Google, and Twitter) that provide budget pacing services based on historical information on all advertisers, may also offer, under conditions that are common on online ad markets, "private" budget-pacing services that do

not use information on competing advertisers, with practically no loss of optimality.

Future directions and open problems. We next list several avenues for future research and discuss limitations and potential extensions. First, our results rely on distribution of values being absolutely continuous and do not easily extend to the case of discrete values. In this case, multiple advertisers might submit the same bid and how ties are broken is important. Conitzer et al. (2017) argue that, on top of shading bids, at equilibrium advertisers employ mixed strategies in which they randomize over bids. We leave the analysis of discrete values as a future research direction.

Our results establish the asymptotic optimality of adaptive pacing strategies in both the “pessimistic” and “optimistic” scenarios when advertisers have no prior knowledge of whether the competitive nature is stationary or arbitrary, and that selecting $\epsilon_k \sim T^{-1/2}$ guarantees a convergence rate of $T^{-1/2}$ in both frameworks. We conjecture that in either of these frameworks, this rate cannot be improved by *any* admissible and budget-feasible strategy (random or not), even when advertisers know up front the environment in which they are competing (i.e., stationary or arbitrary). Kleinberg (2005) provides some evidence that this rate is tight in the stationary setting, showing that no algorithm can achieve a rate better than $B^{-1/2}$ in a variation of the secretary problem in which the decision maker is allowed to choose the best B secretaries. While this result assumes that values are drawn without replacement from an unknown set of numbers (i.e., a random permutation model), we believe that similar results apply to a setting in which values are drawn independently from an unknown distribution.

The update rule of the adaptive pacing strategy takes a simple additive form, but one may consider other update rules. One approach is to use non-linear update rules such as exponentiated gradient descent in which multipliers are updated multiplicatively. One may also use rules that depend on the entire history of observations rather than only on the most recent period, e.g., an update rule in which the gradient step at time t depends on the difference between the target expenditure rate and the average expenditure up to time t , rather than only on the expenditure in period t . Another approach is to use methods that approximate the Hessian of the dual objective function based on all observed gradients such as the online Newton method (Hazan et al. 2007) or the adaptive subgradient method (Duchi et al. 2011).

A Proofs of main results

A.1 Proof of Theorem 3.1

We show that for any $\gamma < \bar{v}_k/\rho_k$ no admissible bidding strategy (including randomized ones) can guarantee asymptotic γ -competitiveness. In the proof we use Yao's Principle, according to which in order to analyze the worst-case performance of randomized algorithms, it suffices to analyze the expected performance of deterministic algorithms given distribution over inputs.

Preliminaries. In §2 we denoted by \mathcal{B}_k the class of admissible bidding strategies defined by the mappings $\{b_{k,t}^\beta : t = 1, \dots, T\}$ together with the distribution \mathbb{P}_y , and the target expenditure rate ρ_k . We now denote by $\tilde{\mathcal{B}}_k \subset \mathcal{B}_k$ the subclass of *deterministic* admissible strategies, defined by adjusting the histories to be $\tilde{\mathcal{H}}_{k,t} := \sigma\left(\left\{v_{k,\tau}, b_{k,\tau}, z_{k,\tau}, u_{k,\tau}\right\}_{\tau=1}^{t-1}, v_{k,t}\right)$ for any $t \geq 2$, with $\tilde{\mathcal{H}}_{k,1} := \sigma(v_{k,1})$. Then, $\tilde{\mathcal{B}}_k$ is the subclass of bidding strategies $\beta \in \mathcal{B}_k$ such that $b_{k,t}^\beta$ is measurable with respect to the filtration $\tilde{\mathcal{H}}_{k,t}$ for each $t \in \{1, \dots, T\}$.

To simplify notation we now drop the dependence on k . Given a target expenditure rate ρ , strategy β , vector of values \mathbf{v} , vector of highest competing bids \mathbf{d} , and a constant $\gamma \geq 1$, we define:

$$R_\gamma^\beta(\mathbf{v}; \mathbf{d}) = \frac{1}{T} \left(\pi^H(\mathbf{v}; \mathbf{d}) - \gamma \pi^\beta(\mathbf{v}; \mathbf{d}) \right).$$

We next adapt Yao's principle (Yao, 1977) to our setting to establish a connection between the expected performance of any randomized strategy and the expected performance of the best deterministic strategy under some distribution over sequences of valuations.

Lemma A.1. (Yao's principle) *Let $\mathbb{E}_{\mathbf{d}}[-]$ denote expectation with respect to some distribution over a set of competing bid sequences $\{\mathbf{d}^1, \dots, \mathbf{d}^m\} \in \mathbb{R}_+^{T \times m}$. Then, for any vector of valuations $\mathbf{v}' \in \mathbb{R}_+^T$ and for any bidding strategy $\beta \in \mathcal{B}$,*

$$\sup_{\substack{\mathbf{v} \in [0, \bar{v}]^T \\ \mathbf{d} \in \mathbb{R}_+^T}} \mathbb{E}^\beta \left[R_\gamma^\beta(\mathbf{v}; \mathbf{d}) \right] \geq \inf_{\beta \in \tilde{\mathcal{B}}} \mathbb{E}_{\mathbf{d}} \left[R_\gamma^\beta(\mathbf{v}'; \mathbf{d}) \right],$$

Lemma A.1 is an adaptation of Theorem 8.3 in Borodin and El-Yaniv (1998); For completeness we provide a proof for this Lemma in Appendix B. By Lemma A.1, to bound the worst-case loss of any admissible strategy (deterministic or not) relative to the best response in hindsight, it suffices to analyze the expected loss of deterministic strategies relative to the same benchmark, where competing bids are drawn from a carefully selected distribution.

Worst-case instance. Fix $T \geq \bar{v}/\rho$ and a target expenditure rate $0 < \rho \leq \bar{v}$. Suppose that the advertiser is a priori informed that the vector of best competitors' bids is

$$\bar{\mathbf{v}} = \left(\underbrace{\bar{v}, \dots, \bar{v}}_{m \lfloor \frac{T}{m} \rfloor \text{ auctions}}, \underbrace{\bar{v}, \bar{v}, \dots, \bar{v}}_{T_0 \text{ auctions}} \right),$$

where we decompose the sequence in $m := \lceil \bar{v}/\rho \rceil$ batches of length $\lfloor T/m \rfloor$, and $T_0 := T - m \lfloor T/m \rfloor$ auctions are added to the end of the sequence to handle cases when T is not divisible by m . Suppose that the advertiser also knows that the sequence of competing bids \mathbf{d} is selected randomly according to a discrete distribution $\mathbf{p} = \{p_1, \dots, p_m\}$ over the set $\mathcal{D} = \{\mathbf{d}^1, \mathbf{d}^2, \dots, \mathbf{d}^m\} \in [0, \bar{v}]^{T \times m}$, where:

$$\begin{aligned} \mathbf{d}^1 &= \left(\underbrace{d_1, \dots, d_1}_{\lfloor \frac{T}{m} \rfloor \text{ auctions}}, \underbrace{d_2, \dots, d_2}_{\lfloor \frac{T}{m} \rfloor \text{ auctions}}, \dots, \underbrace{d_{m-1}, \dots, d_{m-1}}_{\lfloor \frac{T}{m} \rfloor \text{ auctions}}, \underbrace{d_m, \dots, d_m}_{\lfloor \frac{T}{m} \rfloor \text{ auctions}}, \underbrace{\bar{v}, \dots, \bar{v}}_{T_0 \text{ auctions}} \right) \\ \mathbf{d}^2 &= \left(\underbrace{d_1, \dots, d_1}_{\lfloor \frac{T}{m} \rfloor \text{ auctions}}, \underbrace{d_2, \dots, d_2}_{\lfloor \frac{T}{m} \rfloor \text{ auctions}}, \dots, \underbrace{d_{m-1}, \dots, d_{m-1}}_{\lfloor \frac{T}{m} \rfloor \text{ auctions}}, \underbrace{\bar{v}, \dots, \bar{v}}_{\lfloor \frac{T}{m} \rfloor \text{ auctions}}, \underbrace{\bar{v}, \dots, \bar{v}}_{T_0 \text{ auctions}} \right) \\ &\vdots \\ \mathbf{d}^{m-1} &= \left(\underbrace{d_1, \dots, d_1}_{\lfloor \frac{T}{m} \rfloor \text{ auctions}}, \underbrace{d_2, \dots, d_2}_{\lfloor \frac{T}{m} \rfloor \text{ auctions}}, \underbrace{\bar{v}, \dots, \bar{v}}_{(m-2) \lfloor \frac{T}{m} \rfloor \text{ auctions}}, \underbrace{\bar{v}, \bar{v}, \dots, \bar{v}}_{T_0 \text{ auctions}} \right) \\ \mathbf{d}^m &= \left(\underbrace{d_1, \dots, d_1}_{\lfloor \frac{T}{m} \rfloor \text{ auctions}}, \underbrace{\bar{v}, \dots, \bar{v}}_{(m-1) \lfloor \frac{T}{m} \rfloor \text{ auctions}}, \underbrace{\bar{v}, \bar{v}, \dots, \bar{v}}_{T_0 \text{ auctions}} \right), \end{aligned}$$

with $d_j = \bar{v}(1 - \varepsilon^{m-j+1})$ for every $i \in \{1, \dots, m\}$. We use indices $i \in \{1, \dots, m\}$ to refer to competing bid sequences and indices $j \in \{1, \dots, m\}$ to refer to batches within a sequence (other than the last batch which always has zero utility). The parameter ε is such that $\varepsilon \in (0, 1]$; the precise value will be determined later on.

Sequence \mathbf{d}^1 represents a case where competing bids gradually decrease throughout the campaign horizon (except the last T_0 auctions); sequences $\mathbf{d}^2, \dots, \mathbf{d}^m$ follow essentially the same structure, but introduce the risk of net utilities going down to zero at different time points. Thus, the feasible set of competing bids present the advertiser with the following tradeoff: early auctions introduce return per unit of budget that is certain, but low. Later auctions introduce return per unit of budget that may be higher, but uncertain because the net utility (value minus payment) may diminish to zero. In the rest of the proof, the parameters of this instance are tuned to maximize the worst-case loss that must be incurred due to this tradeoff by *any* admissible bidding strategy.

Useful subclass of deterministic strategies. As we next show, under the structure of the worst-case instance described above, it suffices to analyze strategies that determine before the beginning of the campaign how many auctions to win at different stages of the campaign horizon. Define the set $\mathcal{X} := \left\{ \mathbf{x} \in \{0, \dots, \lfloor T/m \rfloor\}^m : \sum_{j=1}^m d_j x_j \leq B \right\}$. Given $\mathbf{x} \in \mathcal{X}$, we denote by $\beta^{\mathbf{x}} \in \tilde{\mathcal{B}}$ a strategy that for each $j \in \{1, \dots, m\}$ bids d_j in the first x_j auctions of batch j . When competing bids are d_j in the j 'th batch, this strategy is guaranteed to win x_j auctions.

Lemma A.2. *For any deterministic bidding strategy $\beta \in \tilde{\mathcal{B}}$ there exists a vector $\mathbf{x} \in \mathcal{X}$ such that for any $\mathbf{d} \in \mathcal{D}$ one has $\pi^{\beta^{\mathbf{x}}}(\bar{\mathbf{v}}, \mathbf{d}) = \pi^{\beta}(\bar{\mathbf{v}}, \mathbf{d})$.*

Lemma A.2 implies that under the structure at hand, the minimal loss that is achievable by a deterministic strategy relative to the hindsight solution is attained within the set \mathcal{X} . The proof of Lemma A.2 follows because every competing bid sequence \mathbf{d}^i is identical to \mathbf{d}^1 , thus indistinguishable, up to time τ^i , the first time at which utilities go down to zero in competing bid sequence \mathbf{d}^i . Therefore, the bids of any deterministic strategy coincide up to time τ^i under competing bid sequences \mathbf{d}^i and \mathbf{d}^1 . Because utilities after time τ^i are zero, all bids under competing bid sequence \mathbf{d}^i after time τ^i are irrelevant. Since all items in a batch have competing bid, it suffices to look at the number of auctions won by the strategy in each batch when the vector of competing bids is \mathbf{d}^1 .

Analysis. Fix $\mathbf{x} \in \mathcal{X}$ and consider the strategy $\beta^{\mathbf{x}}$. Define the net utilities matrix $U \in \mathbb{R}^{m \times m}$:

$$U = \begin{bmatrix} u_1 & u_2 & \dots & u_{m-1} & u_m \\ u_1 & u_2 & \dots & u_{m-1} & 0 \\ \vdots & \vdots & & & \vdots \\ u_1 & u_2 & 0 & \dots & 0 \\ u_1 & 0 & 0 & \dots & 0 \end{bmatrix}$$

with $u_j = \bar{v} - d_j = \bar{v}\varepsilon^{m-j+1}$. The matrix U is invertible with rows representing different sequences in \mathcal{D} and columns capturing the net utility of different batches (perhaps except the last T_0 auctions). Define the vector $\mathbf{u} = (u_m, u_{m-1}, \dots, u_1)' \in \mathbb{R}^m$ as the net utility of the most profitable batch of each sequence in \mathcal{D} . Define the following distribution over sequences of values:

$$\mathbf{p}' = \frac{\mathbf{e}'U^{-1}}{\mathbf{e}'U^{-1}\mathbf{e}},$$

where $\mathbf{e} = (1, \dots, 1)'$ and where U^{-1} denotes the inverse of U . Some algebra shows that

$$\mathbf{e}'U^{-1} = \frac{1}{\bar{v}} \left(\frac{1}{\varepsilon}, \frac{1}{\varepsilon^2} - \frac{1}{\varepsilon}, \frac{1}{\varepsilon^3} - \frac{1}{\varepsilon^2}, \dots, \frac{1}{\varepsilon^m} - \frac{1}{\varepsilon^{m-1}} \right),$$

and $\mathbf{e}'U^{-1}\mathbf{e} = 1/(\bar{v}\varepsilon^m)$. This implies that $p_i \geq 0$ for all $i \in \{1, \dots, m\}$ because $\varepsilon \in (0, 1]$. Since, $\sum_{i=1}^m p_i = 1$ we have that \mathbf{p} is a valid probability distribution over inputs. Then, the expected performance-per-round of β^x is:

$$\frac{1}{T}\mathbb{E}_{\mathbf{d}} \left[\pi^{\beta^x}(\bar{\mathbf{v}}; \mathbf{d}) \right] = \frac{1}{T}\mathbf{p}'U\mathbf{x} = \frac{1}{T}\bar{v}\varepsilon^m\mathbf{e}'\mathbf{x} \stackrel{(a)}{\leq} \varepsilon^m \frac{\rho}{1-\varepsilon}, \quad (\text{A-1})$$

where (a) holds by Lemma A.2 because $B \geq \sum_{j=1}^m d_j x_j \geq d_m \mathbf{e}'\mathbf{x}$ because the competing bids d_j are decreasing and $x_j \geq 0$, and using that $\rho = B/T$ and $d_m = \bar{v}(1-\varepsilon)$. In addition, the expected performance-per-round of the hindsight solution satisfies:

$$\begin{aligned} \frac{1}{T}\mathbb{E}_{\mathbf{d}} \left[\pi^{\text{H}}(\bar{\mathbf{v}}; \mathbf{d}) \right] &\stackrel{(a)}{\geq} \frac{1}{T} \left\lfloor \frac{T}{m} \right\rfloor \mathbf{p}'\mathbf{u} \stackrel{(b)}{=} \frac{1}{T} \left\lfloor \frac{T}{m} \right\rfloor \varepsilon^m (m - \varepsilon(m-1)) \\ &> \left(\frac{1}{m} - \frac{1}{T} \right) \bar{v}\varepsilon^m (m - \varepsilon(m-1)), \end{aligned} \quad (\text{A-2})$$

where: (a) holds since by the definition of the index m one has that $\rho \geq \bar{v}/m$ and therefore $B = T\rho \geq T\bar{v}/m \geq \lfloor T/m \rfloor d_j$ because $d_j \leq \bar{v}$, so the hindsight solution allows winning at least the $\lfloor T/m \rfloor$ most profitable auctions; and (b) holds since $\mathbf{p}'\mathbf{u} = (\mathbf{e}'U^{-1}\mathbf{u})/(\mathbf{e}'U^{-1}\mathbf{e}) = \bar{v}\varepsilon^m(m-\varepsilon(m-1))$.

Set $\delta \in (0, 1)$ such that $\gamma = (1-\delta)(\bar{v}/\rho)$. Setting $\varepsilon = \frac{\delta}{\delta+4} \leq 1$, one has for any $T \geq 4 \lceil \bar{v}/\rho \rceil / \delta$:

$$\begin{aligned} \mathbb{E}_{\mathbf{d}} \left[R_{\gamma}^{\beta^x}(\bar{\mathbf{v}}; \mathbf{d}) \right] &\stackrel{(a)}{\geq} \bar{v}\varepsilon^m \left(\left(\frac{1}{m} - \frac{1}{T} \right) (m - \varepsilon(m-1)) - \frac{1-\delta}{1-\varepsilon} \right) \\ &= \bar{v}\varepsilon^m \left(\frac{\delta}{1-\varepsilon} - \varepsilon \left(\frac{m-1}{m} + \frac{1}{1-\varepsilon} \right) - \frac{m}{T} + \varepsilon \frac{m-1}{T} \right) \\ &\stackrel{(b)}{\geq} \bar{v}\varepsilon^m \left(\delta - \frac{2\varepsilon}{1-\varepsilon} - \frac{m}{T} \right) \\ &\stackrel{(c)}{=} \bar{v} \left(\frac{\delta}{\delta+4} \right)^m \left(\frac{\delta}{2} - \frac{m}{T} \right) \\ &\stackrel{(d)}{\geq} \frac{\bar{v}}{4} \left(\frac{1}{\delta+4} \right)^m \delta^{m+1} > 0, \end{aligned}$$

where: (a) holds from (A-1) and (A-2); (b) holds by using that $1/(1-\varepsilon) \geq 1$ because $0 \leq \varepsilon \leq \bar{1}$ for the first term in the parenthesis, using that $(m-1)/m \leq 1 \leq 1/(1-\varepsilon)$ in the second term, and dropping the last term because $m \geq 2$; (c) follows by setting $\varepsilon = \frac{\delta}{\delta+4}$; and (d) holds because $\delta/2 - m/T \geq \delta/4$ for any $T \geq 4 \lceil \bar{v}/\rho \rceil / \delta$ since $m = \lceil \frac{\bar{v}}{\rho} \rceil$ and $\delta \in (0, 1)$. This establishes that for

any $\gamma < \bar{v}/\rho$ there is a constant $\underline{C} > 0$ such that for T large enough:

$$\inf_{\beta \in \mathcal{B}} \sup_{\substack{\mathbf{v} \in [0, \bar{v}]^T \\ \mathbf{d} \in \mathbb{R}_+^T}} \mathbb{E}^\beta \left[R_\gamma^\beta(\mathbf{v}; \mathbf{d}) \right] \stackrel{(a)}{\geq} \inf_{\beta \in \tilde{\mathcal{B}}} \mathbb{E}_{\mathbf{d}} \left[R_\gamma^\beta(\bar{\mathbf{v}}; \mathbf{d}) \right] \stackrel{(b)}{\geq} \inf_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{\mathbf{d}} \left[R_\gamma^{\beta^x}(\bar{\mathbf{v}}; \mathbf{d}) \right] \geq \underline{C},$$

where (a) follows from Lemma A.1, and (b) holds by Lemma A.2. Therefore, no admissible bidding strategy can guarantee γ -competitiveness for any $\gamma < \bar{v}/\rho$, concluding the proof. \square

A.2 Proof of Theorem 3.3

To simplify notation we drop the dependence on k . We consider an alternate process in which multipliers and bids continue being updated after the budget is depleted, but no reward is accrued after budget depletion. Denote by $\tilde{\tau}^\Lambda := \inf\{t \geq 1 : \tilde{B}_t^\Lambda < \bar{v}\}$ the first auction in which the bid of the advertiser is constrained by the remaining budget, and by $\tau^\Lambda := \tilde{\tau}^\Lambda - 1$ the last period in which bids are unconstrained. The dynamics of the adaptive pacing strategy are:

$$\begin{aligned} \mu_{t+1} &= P_{[0, \bar{\mu}]}(\mu_t - \epsilon(\rho - d_t x_t^\Lambda)) \\ x_t^\Lambda &= \mathbf{1}\{v_t - d_t \geq \mu_t d_t\} \\ \tilde{B}_{t+1}^\Lambda &= \tilde{B}_t^\Lambda - d_t x_t^\Lambda. \end{aligned}$$

Given the sequence of competing bids \mathbf{d} we denote by $z_t^\Lambda = d_t x_t^\Lambda$ the expenditure at time t under the adaptive pacing strategy. Because the auction is won only if $v_t - d_t \geq \mu_t d_t$, the expenditure satisfies $z_t^\Lambda \leq v_t/(1 + \mu_t) \leq \bar{v}$.

Step 1: controlling the stopping time $\tilde{\tau}^\Lambda$. We first show that the adaptive pacing strategy does not run out of budget too early, that is, τ^Λ is close to T . From the update rule of the dual variables one has for every $t \leq \tau^\Lambda$,

$$\mu_{t+1} = P_{[0, \bar{\mu}]}(\mu_t + \epsilon(z_t^\Lambda - \rho)) = \mu_t + \epsilon(z_t^\Lambda - \rho) - P_t,$$

where we define $P_t := \mu_t + \epsilon(z_t^\Lambda - \rho) - P_{[0, \bar{\mu}]}(\mu_t + \epsilon(z_t^\Lambda - \rho))$ as the projection error. Reordering terms and summing up to period τ^Λ one has

$$\sum_{t=1}^{\tau^\Lambda} (z_t^\Lambda - \rho) = \sum_{t=1}^{\tau^\Lambda} \frac{1}{\epsilon} (\mu_{t+1} - \mu_t) + \sum_{t=1}^{\tau^\Lambda} \frac{P_t}{\epsilon}. \quad (\text{A-3})$$

We next bound each term independently. For left-hand side of (A-3) we have

$$\sum_{t=1}^{\tau^A} (z_t^A - \rho) \stackrel{(a)}{=} B - \tilde{B}_{\tau^A+1}^A - \rho\tau^A \stackrel{(b)}{\geq} \rho(T - \tau^A) - \bar{v},$$

where (a) holds since $\tilde{B}_{\tau^A+1}^A = B - \sum_{t=1}^{\tau^A} z_t^A$ and (b) uses $\tilde{B}_{\tau^A+1}^A \leq \bar{v}$ and $\rho = B/T$. For the first term in the right-hand side of (A-3) we have

$$\sum_{t=1}^{\tau^A} \frac{1}{\epsilon} (\mu_{t+1} - \mu_t) = \frac{\mu_{\tau^A+1}^A}{\epsilon} - \frac{\mu_1^A}{\epsilon} \leq \frac{\bar{\mu}}{\epsilon},$$

where the inequality follows because $\mu_t \in [0, \bar{\mu}_k]$. We next bound the second term in the right-hand side of (A-3). The projection error satisfies

$$\begin{aligned} P_t &\leq P_t^+ = (\mu_t + \epsilon(z_t^A - \rho) - P_{[0, \bar{\mu}]}(\mu_t + \epsilon(z_t^A - \rho)))^+ \\ &\stackrel{(a)}{=} (\mu_t + \epsilon(z_t^A - \rho) - \bar{\mu}) \mathbf{1}\{\mu_t + \epsilon(z_t^A - \rho) > \bar{\mu}\} \\ &\stackrel{(b)}{\leq} \epsilon \bar{v} \mathbf{1}\{\mu_t + \epsilon(z_t^A - \rho) > \bar{\mu}\} \end{aligned}$$

where (a) holds since the projection error is positive only if $\mu_t + \epsilon(z_t^A - \rho) > \bar{\mu}$, and (b) holds since $\mu_t \leq \bar{\mu}$ and since the deviation is bounded by $z_t^A - \rho \leq z_t^A \leq \bar{v}$ as the expenditure is at most \bar{v} . We next show that there is no positive projection error, or more formally, $P_t^+ = 0$ whenever $\epsilon\bar{v} \leq 1$. Consider the function $f: \mathbb{R}_+ \rightarrow \mathbb{R}$ given by $f(\mu) = \mu + (\epsilon\bar{v})/(1 + \mu)$. The function $f(\cdot)$ is non-decreasing whenever $\epsilon\bar{v} \leq 1$, and therefore

$$\begin{aligned} \mu_t + \epsilon(z_t^A - \rho) &\stackrel{(a)}{\leq} \mu_t + \epsilon \left(\frac{\bar{v}}{1 + \mu_t} - \rho \right) = f(\mu_t) - \epsilon\rho \\ &\stackrel{(b)}{\leq} f(\bar{\mu}) - \frac{\epsilon\bar{v}}{1 + \bar{\mu}} \stackrel{(c)}{=} \bar{\mu}, \end{aligned}$$

where (a) holds by $z_t^A = d_t \mathbf{1}\{v_t - d_t \geq \mu_t d_t\} \leq v_t/(1 + \mu_t) \leq \bar{v}/(1 + \mu_t)$ as a payment is never greater than $v_t/(1 + \mu_t)$, and (b) holds since $f(\cdot)$ is non-decreasing, $\mu_t \leq \bar{\mu}$, and $\bar{\mu} \geq \bar{v}/\rho - 1$. Therefore, there is no positive projection error when $\epsilon\bar{v} \leq 1$. Using these inequalities in (A-3) one has

$$T - \tau^A \leq \frac{\bar{\mu}}{\epsilon\rho} + \frac{\bar{v}}{\rho}. \tag{A-4}$$

Let $x \wedge y = \min(x, y)$. Therefore the strategy does not run out of budget very early, and one has

$$\begin{aligned}
\pi^{\text{A}}(\mathbf{v}; \mathbf{d}) &\geq \sum_{t=1}^{T \wedge \tau^{\text{A}}} (v_t - d_t) x_t^{\text{A}} \\
&\stackrel{(a)}{\geq} \sum_{t=1}^T (v_t - d_t) x_t^{\text{A}} - \bar{v} (T - \tau^{\text{A}})^+ \\
&\stackrel{(b)}{\geq} \sum_{t=1}^T (v_t - d_t) x_t^{\text{A}} - \underbrace{\left(\frac{\bar{v} \bar{\mu}}{\epsilon \rho} + \frac{\bar{v}^2}{\rho} \right)}_{E_1}, \tag{A-5}
\end{aligned}$$

where (a) follows from $0 \leq v_t - d_t \leq \bar{v}$, and (b) follows from (A-4).

Step 2: bounding the regret. We next bound the relative loss in terms of the potential value of auctions lost by the adaptive pacing strategy. Consider a relaxation of the hindsight problem (1) in which the decision maker is allowed to take fractional items. In this relaxation, the optimal solution is a greedy strategy that sorts items in decreasing order of the ratio $(v_t - d_t)/d_t$ and wins items until the budget is depleted (where the “overflowing” item is fractionally taken). Let μ^{H} be the sample-path dependent fixed threshold corresponding to the ratio of the “overflowing” item. Then all auctions won satisfy $(v_t - d_t) \geq \mu^{\text{H}} d_t$. We obtain an upper bound by assuming that the overflowing item is completely taken by the hindsight solution. The dynamics of this bound are:

$$\begin{aligned}
x_t^{\text{H}} &= \mathbf{1}\{v_t - d_t \geq \mu^{\text{H}} d_t\} \\
\tilde{B}_{t+1}^{\text{H}} &= \tilde{B}_t^{\text{H}} - x_t^{\text{H}} z_t^{\text{H}}.
\end{aligned}$$

Denote by $\tilde{\tau}^{\text{H}} := \inf\{t \geq 1 : \tilde{B}_t^{\text{H}} < 0\}$ the period in which the budget of the hindsight solution is depleted, and by $\tau^{\text{H}} := \tilde{\tau}^{\text{H}} - 1$. Then, one has:

$$\begin{aligned}
\pi^{\text{H}}(\mathbf{v}; \mathbf{d}) - \pi^{\text{A}}(\mathbf{v}; \mathbf{d}) &\leq \sum_{t=1}^{T \wedge \tau^{\text{H}}} x_t^{\text{H}} (v_t - d_t) - \sum_{t=1}^{T \wedge \tau^{\text{A}}} x_t^{\text{A}} (v_t - d_t) \\
&\stackrel{(a)}{\leq} \sum_{t=1}^{T \wedge \tau^{\text{H}}} (v_t - d_t) (x_t^{\text{H}} - x_t^{\text{A}}) + E_1 \tag{A-6}
\end{aligned}$$

where (a) follow from (A-5), ignoring periods after τ^{H} . In addition, for any $t \in \{1, \dots, T \wedge \tau^{\text{H}}\}$:

$$\begin{aligned}
(v_t - d_t) (x_t^H - x_t^A) &\stackrel{(a)}{=} (v_t - d_t) (\mathbf{1}\{v_t - d_t \geq \mu^H d_t\} - \mathbf{1}\{v_t - d_t \geq \mu_t d_t\}) \\
&\stackrel{(b)}{=} (v_t - d_t) (\mathbf{1}\{\mu_t d_t > v_t - d_t \geq \mu^H d_t\} - \mathbf{1}\{\mu^H d_t > v_t - d_t \geq \mu_t d_t\}) \\
&\stackrel{(c)}{\leq} (v_t - d_t) \mathbf{1}\{\mu_t d_t > v_t - d_t\} = (v_t - d_t) (1 - x_t^A), \tag{A-7}
\end{aligned}$$

where: (a) follows the threshold structure of the hindsight solution as well as the dynamics of the adaptive pacing strategy; (b) holds since $x_t^H - x_t^A = 1$ if and only if the auction is won in the hindsight solution but not by the adaptive pacing strategy, and $x_t^H - x_t^A = -1$ if and only if the auction is won by the adaptive pacing strategy but not in the hindsight solution; and (c) follows from discarding indicators. Putting (A-6) and (A-7) together one obtains a bound on the relative loss in terms of the potential value of the auctions lost by the adaptive pacing strategy:

$$\pi^H(\mathbf{v}; \mathbf{d}) - \pi^A(\mathbf{v}; \mathbf{d}) \leq \sum_{t=1}^{T \wedge \tau^H} (v_t - d_t) (1 - x_t^A) + E_1. \tag{A-8}$$

Step 3: bounding the potential value of lost auctions. We next bound the value of the auctions lost under the adaptive pacing strategy in terms of the value of alternative auctions won under the same strategy. We show that for all $t = 1, \dots, T$:

$$(v_t - d_t) (1 - x_t^A) \leq \left(\frac{\bar{v}}{\rho} - 1\right) (v_t - d_t) x_t^A + \frac{\bar{v}}{\rho} \mu_t (\rho - z_t^A). \tag{A-9}$$

Note that when $d_t > \bar{v}$, (A-9) simplifies to $v_t - d_t \leq \bar{v} \mu_t$, because $x_t^A = 0$ and $z_t^A = d_t x_t^A = 0$. The inequality then holds since $v_t - d_t < 0$ and $\mu_t \geq 0$. We next prove (A-9) holds also when $d_t \leq \bar{v}$. Using $z_t^A = d_t \mathbf{1}\{v_t - d_t \geq \mu_t d_t\}$ one obtains

$$\begin{aligned}
\mu_t (\rho - z_t^A) &\stackrel{(a)}{\geq} \mu_t \left(\frac{\rho}{\bar{v}} d_t - z_t^A\right) = \mu_t d_t \left(\frac{\rho}{\bar{v}} - \mathbf{1}\{v_t - d_t \geq \mu_t d_t\}\right) \\
&= \frac{\rho}{\bar{v}} \mu_t d_t \mathbf{1}\{\mu_t d_t > v_t - d_t\} - \left(1 - \frac{\rho}{\bar{v}}\right) \mu_t d_t \mathbf{1}\{v_t - d_t \geq \mu_t d_t\} \\
&\stackrel{(b)}{\geq} \frac{\rho}{\bar{v}} (v_t - d_t) (1 - x_t^A) - \left(1 - \frac{\rho}{\bar{v}}\right) (v_t - d_t) x_t^A,
\end{aligned}$$

where (a) holds because $d_t \leq \bar{v}$, and (b) holds because $\rho \leq \bar{v}$ and using $\mu_t d_t > v_t - d_t$ in the first term and $\mu_t d_t \leq v_t - d_t$ in the second term. The claim follows from multiplying by \bar{v}/ρ and

reordering terms. Summing (A-9) over $t = 1 \dots, T$ implies

$$\sum_{t=1}^T (v_t - d_t) (1 - x_t^A) \leq \left(\frac{\bar{v}}{\rho} - 1 \right) \sum_{t=1}^T (v_t - d_t) x_t^A + \frac{\bar{v}}{\rho} \sum_{t=1}^T \mu_t (\rho - z_t^A). \quad (\text{A-10})$$

We next bound the second term in (A-10). The update rule of the strategy implies that for any $\mu \in [0, \bar{\mu}]$ one has

$$\begin{aligned} (\mu_{t+1} - \mu)^2 &\stackrel{(a)}{\leq} (\mu_t - \mu - \epsilon(\rho - z_t^A))^2 \\ &= (\mu_t - \mu)^2 - 2\epsilon(\mu_t - \mu)(\rho - z_t^A) + \epsilon^2(\rho - z_t^A)^2. \end{aligned}$$

where (a) follows from a standard contraction property of the Euclidean projection operator. Re-ordering terms and summing over $t = 1, \dots, T$ we have for $\mu = 0$:

$$\begin{aligned} \sum_{t=1}^T \mu_t (\rho - z_t^A) &\leq \sum_{t=1}^T \frac{1}{2\epsilon} \left((\mu_t)^2 - (\mu_{t+1})^2 \right) + \frac{\epsilon}{2} (\rho - z_t^A)^2 \\ &\stackrel{(a)}{=} \frac{(\mu_1^A)^2}{2\epsilon} - \frac{(\mu_{T+1})^2}{2\epsilon} + \sum_{t=1}^T \frac{\epsilon}{2} (\rho - z_t^A)^2 \\ &\stackrel{(b)}{\leq} \frac{\bar{\mu}^2}{2\epsilon} + \frac{\bar{v}^2}{2} T\epsilon, \end{aligned} \quad (\text{A-11})$$

where (a) follows from telescoping the sum, and (b) follows from $\mu_t \in [0, \bar{\mu}]$ together with $\rho, z_t^A \in [0, \bar{v}]$. Combining (A-10) and (A-11) implies

$$\sum_{t=1}^T (v_t - d_t) (1 - x_t^A) \leq \left(\frac{\bar{v}}{\rho} - 1 \right) \sum_{t=1}^T (v_t - d_t) x_t^A + \underbrace{\left(\frac{\bar{v}\bar{\mu}^2}{2\rho\epsilon} + \frac{\bar{v}^3}{2\rho} T\epsilon \right)}_{E_2}. \quad (\text{A-12})$$

Step 4: putting everything together. Using the bound on regret derived in (A-8) we have:

$$\begin{aligned}
\pi^{\text{H}}(\mathbf{v}; \mathbf{d}) - \pi^{\text{A}}(\mathbf{v}; \mathbf{d}) &\leq \sum_{t=1}^{T \wedge \tau^{\text{H}}} (v_t - d_t) (1 - x_t^{\text{A}}) + E_1 \\
&\leq \sum_{t=1}^T (v_t - d_t) (1 - x_t^{\text{A}}) + E_1 \\
&\stackrel{(b)}{\leq} \left(\frac{\bar{v}}{\rho} - 1 \right) \sum_{t=1}^T (v_t - d_t) x_t^{\text{A}} + E_1 + E_2 \\
&\stackrel{(c)}{\leq} \left(\frac{\bar{v}}{\rho} - 1 \right) \pi^{\text{A}}(\mathbf{v}; \mathbf{d}) + \frac{\bar{v}}{\rho} E_1 + E_2,
\end{aligned}$$

where: (a) follows from adding all (non-negative) terms after τ^{H} and since $d_t \leq v_t$ for all $t = 1, \dots, T$; (b) follows from (A-12); and (c) follows from (A-5). Reordering terms we obtain that

$$\pi^{\text{H}}(\mathbf{v}; \mathbf{d}) - \frac{\bar{v}}{\rho} \pi^{\text{A}}(\mathbf{v}; \mathbf{d}) \leq \frac{\bar{v}}{\rho} E_1 + E_2 = \frac{\bar{v}^3}{\rho^2} + \left(\frac{\bar{v}^2 \bar{\mu}}{\rho^2} + \frac{\bar{v} \bar{\mu}^2}{2\rho} \right) \frac{1}{\epsilon} + \frac{\bar{v}^3}{2\rho} T \epsilon,$$

and the result follows. \square

A.3 Proof of Theorem 3.4

To simplify notation we drop the dependence on k . We prove the result in five steps. We first upper bound the expected performance in hindsight using the optimal dual objective value $\Psi(\mu^*)$. Then, we lower bound the performance of an adaptive pacing strategy by discarding the utility of all auctions after budget depletion. We perform a second-order expansion of the expected utility per auction around μ^* to obtain a lower quadratic envelope involving $\Psi(\mu^*)/T$ as the zeroth-order term, the absolute mean error as the first-order term, and the mean squared error as the second-order term. The next step involves upper bounding the absolute mean errors and the mean squared errors incurred throughout the horizon. We conclude by combining these steps and using that the time of budget depletion under the strategy is “close” to T .

Step 1: upper bound on the performance in hindsight. Taking expectations in equation (4) and using Jensen’s inequality together values and competing bids being i.i.d, one obtains that

$$\mathbb{E}_{\mathbf{v}, \mathbf{d}} [\pi^{\text{H}}(\mathbf{v}; \mathbf{d})] \leq \inf_{\mu \geq 0} \Psi(\mu). \tag{A-13}$$

In the remainder of the proof, we lower bound $\mathbb{E}_{\mathbf{v}, \mathbf{d}} [\pi^{\text{A}}(\mathbf{v}; \mathbf{d})]$ in terms of $\Psi(\mu^*)$, where μ^* minimizes the dual function $\Psi(\mu) = T \left(\mathbb{E}_{v, d} [(v - (1 + \mu)d)^+] + \mu \rho \right)$. Since the dual function is differentiable, the Karush-Kuhn-Tucker optimality condition implies that μ^* satisfies the complementary condition

$$\mu^* \geq 0 \quad \perp \quad G(\mu^*) \leq \rho, \quad (\text{A-14})$$

where \perp denotes a complementarity condition between the multiplier and the expenditure, that is, at least one condition must hold with equality. Here we used that $\bar{\Psi}'(\mu) = \rho - G(\mu)$, where $\bar{\Psi}(\mu) = \Psi(\mu)/T$ and the expenditure function is $G(\mu) := \mathbb{E}_{v,d}[\mathbf{1}\{(1 + \mu)d \leq v\}d]$.

Step 2: lower bound on adaptive pacing's performance. We next lower bound the performance of the strategy. We consider an alternate framework in which the advertiser is allowed to bid even after budget depletion. Let $\tilde{B}_t = B - \sum_{s=1}^{t-1} z_s$ denote the advertiser's remaining budget at the beginning of period t in the alternate framework. Let $\tau = \inf\{t \geq 1 : \tilde{B}_{t+1} < \bar{v}\}$ be the last period in which the remaining budget is larger than \bar{v} . Since $v/(1 + \mu) \leq \bar{v}$ for any $v \in [0, \bar{v}]$ and $\mu \geq 0$, for any period $t \leq \tau$ the bids of the advertiser are $b_t^\Lambda = v_t/(1 + \mu_t)$. Therefore the performance in both the original and alternate frameworks coincide up to time τ , and therefore

$$\mathbb{E}_{\mathbf{v}, \mathbf{d}}[\pi^\Lambda(\mathbf{v}, \mathbf{d})] \stackrel{(a)}{\geq} \mathbb{E}\left[\sum_{t=1}^{\tau \wedge T} u_t\right] \stackrel{(b)}{\geq} \mathbb{E}\left[\sum_{t=1}^T u_t\right] - \bar{v}\mathbb{E}[(T - \tau)^+], \quad (\text{A-15})$$

where (a) follows from discarding all auctions after the time the advertiser runs out of budget; and (b) follows from $0 \leq u_t \leq \bar{v}$. The second term can be bounded by (A-4) because $\epsilon \bar{v} \leq 1$ from Assumption 3.2 and $\bar{\mu} \geq \bar{v}/\rho - 1$. We next bound the first term.

Step 3: lower bound on utility-per-auction. A key step in the proof is to show that the utility per auction collected by the advertiser is “close” to $\bar{\Psi}(\mu^*)$. Using the structure of the strategy, the utility of advertiser k from the t^{th} auction can be written as

$$u_t = \mathbf{1}\{d_t(1 + \mu_t) \leq v_t\}(v_t - d_t) = (v_t - (1 + \mu_t)d_t)^+ + \rho\mu_t + \mu_t(z_t - \rho).$$

Taking expectations conditional on the current multiplier μ_t we obtain

$$\mathbb{E}[u_t \mid \mu_t] = \bar{\Psi}(\mu_t) + \mu_t(G(\mu_t) - \rho),$$

where the equality follows from the linearity of expectation and since v_t and d_t are independent of the multiplier μ_t . Let $U(\mu) = \bar{\Psi}(\mu) + \mu(G(\mu) - \rho)$ be the expected utility per auction when the multiplier is μ . From the complementary slackness conditions (A-14) we have that $U(\mu^*) = \bar{\Psi}(\mu^*)$. Thrice differentiability of the dual function implies that $U(\mu)$ is twice differentiable with derivatives $U'(\mu) = \mu G'(\mu)$ and $U''(\mu) = \mu G''(\mu) + G'(\mu)$ because $\bar{\Psi}'(\mu) = \rho - G(\mu)$ and the product and

summation of differentiable functions is differentiable. Moreover, because the derivatives of $\bar{\Psi}(\mu)$ are bounded, there exists $\bar{G}' > 0$ and $\bar{G}'' > 0$ such that $|G'(\mu)| \leq \bar{G}'$ and $|G''(\mu)| \leq \bar{G}''$, respectively.

Performing a second-order expansion around μ^* and taking expectations we obtain that

$$\begin{aligned} \mathbb{E}[u_t] &= \bar{\Psi}(\mu^*) + \mu^* G'(\mu^*) \mathbb{E}[\mu_t - \mu^*] + \mathbb{E}\left[\frac{U''(\xi)}{2}(\mu_t - \mu^*)^2\right] \\ &\geq \bar{\Psi}(\mu^*) - \underbrace{\bar{G}' \mu^* |\mathbb{E}[\mu_t - \mu^*]|}_{r_t} - \underbrace{\frac{\bar{\mu} \bar{G}'' + \bar{G}'}{2} \mathbb{E}[(\mu_t - \mu^*)^2]}_{\delta_t}, \end{aligned} \quad (\text{A-16})$$

where the first equation follows from Taylor's Theorem for some ξ between μ_t and μ^* , and the inequality follows from the bounds on the derivatives of the expenditure function. We next upper bound the mean squared error δ_t and the absolute mean error r_t .

Step 4: bound on total expected errors. Let $\delta_t = \mathbb{E}[(\mu_t - \mu^*)^2]$ be the mean squared error. Note that the expenditure function G is λ -strongly monotone since the dual function is strongly convex with parameter λ . Because the step size satisfies $2\lambda\epsilon \leq 1$, equation (A-21) implies that

$$\delta_t \leq \bar{\mu}^2 (1 - 2\lambda\epsilon)^{t-1} + \frac{\bar{v}^2}{2\lambda} \epsilon.$$

Thus, the total mean squared error is bounded by

$$\sum_{t=1}^T \delta_t \leq \frac{\bar{v}^2 T}{2\lambda} \epsilon + \bar{\mu}^2 \sum_{t=1}^T (1 - 2\lambda\epsilon)^{t-1} \leq \frac{\bar{v}^2}{2\lambda} T \epsilon + \frac{\bar{\mu}^2}{2\lambda} \frac{1}{\epsilon}, \quad (\text{A-17})$$

where the last inequality follows from $\sum_{t=1}^T (1 - 2\lambda\epsilon)^{t-1} \leq \sum_{t=0}^{\infty} (1 - 2\lambda\epsilon)^t = \frac{1}{2\lambda\epsilon}$.

Let $r_t = \mu^* |\mathbb{E}[\mu_{t+1} - \mu^*]|$ be the absolute mean error (times the optimal multiplier). We proceed under the assumption that $\mu^* > \epsilon\rho$. Otherwise, we can use bound $r_t \leq \epsilon\rho\bar{\mu}$, which is sufficient for our results. From the update rule of the dual variables one has

$$\mu_{t+1} = P_{[0, \bar{\mu}]}(\mu_t + \epsilon(z_t - \rho)) = \mu_t + \epsilon(z_t - \rho) - P_t,$$

where we define $P_t := \mu_t + \epsilon(z_t - \rho) - P_{[0, \bar{\mu}]}(\mu_t + \epsilon(z_t - \rho))$ as the projection error. Subtracting μ^* , multiplying by μ^* and taking expectations we obtain that

$$\begin{aligned} \mu^* \mathbb{E}[\mu_{t+1} - \mu^*] &\stackrel{(a)}{=} \mu^* \mathbb{E}[\mu_t - \mu^*] + \epsilon \mu^* (\mathbb{E}[G(\mu_t)] - G(\mu^*)) - \mu^* \mathbb{E}[P_t] \\ &\stackrel{(b)}{=} (1 + \epsilon G'(\mu^*)) \mu^* \mathbb{E}[\mu_t - \mu^*] + \frac{G''(\xi)}{2} \epsilon \mu^* \mathbb{E}[(\mu_t - \mu^*)^2] - \mu^* \mathbb{E}[P_t], \end{aligned} \quad (\text{A-18})$$

where (a) follows because values v_t and competing bids d_t are independent of the multiplier μ_t , and $\mu^* (G(\mu^*) - \rho) = 0$ from the complementary slackness conditions (A-14); and (b) follows from Taylor's Theorem for some ξ between μ^* and μ_t because G is twice differentiable. We next bound the third term in the right-hand side of (A-18). The projection error satisfies

$$P_t \stackrel{(a)}{\geq} (\mu_t + \epsilon(z_t - \rho)) \mathbf{1} \{ \mu_t + \epsilon(z_t - \rho) < 0 \} \stackrel{(b)}{\geq} -\epsilon\rho \mathbf{1} \{ \mu_t < \epsilon\rho \} ,$$

where (a) holds as the projection error is negative only if $\mu_t + \epsilon(z_t - \rho) < 0$, and (b) holds by $\mu_t \geq 0$ and the expenditure being non-negative. Taking expectations, one has by Markov's inequality that

$$-\mathbb{E}[P_t] \leq \epsilon\rho \mathbb{P} \{ \mu_t < \epsilon\rho \} \leq \epsilon\rho \mathbb{P} \{ |\mu_t - \mu^*| \geq |\mu^* - \epsilon\rho| \} \leq \frac{\epsilon\rho\delta_t}{(\mu^* - \epsilon\rho)^2} .$$

Taking absolute values in (A-18) and invoking the triangle inequality we conclude that

$$r_{t+1} \leq (1 - \lambda\epsilon)r_t + \epsilon \underbrace{\left(\frac{\bar{\mu}\rho}{(\mu^* - \epsilon\rho)^2} + \frac{\bar{\mu}\bar{G}''}{2} \right)}_R \delta_t ,$$

where we used that $\mu^* \leq \bar{\mu}$, $\sup_{\mu \in [0, \bar{\mu}]} G'(\mu) \leq -\lambda$, and $\sup_{\mu \in [0, \bar{\mu}]} G''(\mu) \leq \bar{G}''$ as the dual function has bounded derivatives. By Lemma C.4 with $a = \lambda\epsilon \leq 1$ and $b_t = \epsilon R \delta_t$ together with $\mu^* \leq \bar{\mu}$ and $r_1 \leq \bar{\mu}^2$, one has

$$r_t \leq \bar{\mu}^2 (1 - \lambda\epsilon)^{t-1} + \epsilon R \sum_{s=1}^{t-1} (1 - \lambda\epsilon)^{t-1-s} \delta_s .$$

Thus, the total absolute mean error is bounded by

$$\begin{aligned} \sum_{t=1}^T r_t &\leq \bar{\mu}^2 \sum_{t=1}^T (1 - \lambda\epsilon)^{t-1} + \epsilon R \sum_{t=1}^T \sum_{s=1}^{t-1} (1 - \lambda\epsilon)^{t-1-s} \delta_s \\ &= \bar{\mu}^2 \sum_{t=0}^{T-1} (1 - \lambda\epsilon)^t + \epsilon R \sum_{s=1}^{T-1} \delta_s \sum_{t=0}^{T-1-s} (1 - \lambda\epsilon)^t \\ &\leq \frac{\bar{\mu}^2}{\lambda\epsilon} + \frac{R}{\lambda} \sum_{s=1}^T \delta_s , \end{aligned} \tag{A-19}$$

where the equality follows from changing the order of the summation and the last inequality because $\sum_{t=0}^s (1 - \lambda\epsilon)^t \leq \frac{1}{\lambda\epsilon}$ for all $s \geq 0$.

Step 5: putting everything together. Combining equations (A-13), (A-15), (A-4), summing (A-16) over $t = 1, \dots, T$, and using (A-17) and (A-19) we conclude that there exist constants

$C_1, C_2, C_3 > 0$ such that

$$\mathbb{E}_{\mathbf{v}, \mathbf{d}} [\pi^{\text{H}}(\mathbf{v}; \mathbf{d}) - \pi^{\text{A}}(\mathbf{v}, \mathbf{d})] \leq C_1 + \frac{C_2}{\epsilon} + C_3 T \epsilon.$$

Thus the regret converges to zero under Assumption 3.2, and setting the step size to $\epsilon \sim T^{-1/2}$ yields a convergence rate of $T^{-1/2}$. This concludes the proof. \square

A.4 Proof of Theorem 4.3

We prove the result by first bounding the mean squared error when all bids are unconstrained, i.e., $\tilde{B}_{k,t+1} \geq \bar{v}_k$ for every advertiser k . We then argue that budgets are not depleted too early when the strategy is followed by all advertisers. We conclude by combining these results to upper bound the time average mean squared error.

Step 1. At a high level, this step adapts standard stochastic approximation results (see, e.g., Nemirovski et al. 2009) to the expenditure observations and to accommodate the possibility of different step sizes across agents. Fix some $k \in \{1, \dots, K\}$. The squared error satisfies the recursion

$$\begin{aligned} |\mu_{k,t+1} - \mu_k^*|^2 &= |P_{[0, \bar{\mu}_k]}(\mu_{k,t} - \epsilon_k(\rho_k - z_{k,t})) - \mu_k^*|^2 \\ &\stackrel{(a)}{\leq} |\mu_{k,t} - \mu_k^* - \epsilon_k(\rho_k - z_{k,t})|^2 \\ &= |\mu_{k,t} - \mu_k^*|^2 - 2\epsilon_k(\mu_{k,t} - \mu_k^*)(\rho_k - z_{k,t}) + \epsilon_k^2 |\rho_k - z_{k,t}|^2, \end{aligned}$$

where (a) follows from a standard contraction property of the Euclidean projection operator.

Given a vector of multipliers $\boldsymbol{\mu}$, let $G_k(\boldsymbol{\mu}) := \mathbb{E}_{\mathbf{v}} [\mathbf{1}\{(1 + \mu_k)d_k \leq v_k\}d_{k,1}]$ be the expected expenditure under the second-price auction allocation rule, with $d_k = \max_{i \neq k} \{v_i / (1 + \mu_i)\}$ and the expectation taken w.r.t. the values $v_k \sim F_k$. Define $\delta_{k,t} := \mathbb{E} [(\mu_{k,t} - \mu_k^*)^2 \mathbf{1}\{\tilde{B}_{k,t+1} \geq \bar{v}_k \forall k\}]$ and $\delta_t := \sum_{k=1}^K \delta_{k,t}$. Similarly, define $\hat{\delta}_{k,t} := \delta_{k,t} / \epsilon_k$ and $\hat{\delta}_t := \sum_{k=1}^K \hat{\delta}_{k,t}$. Taking expectations and dividing by ϵ_k we obtain that

$$\begin{aligned} \hat{\delta}_{k,t+1} &\stackrel{(a)}{\leq} \hat{\delta}_{k,t} - 2\mathbb{E} [(\mu_{k,t} - \mu_k^*)(\rho_k - z_{k,t})] + \epsilon_k \mathbb{E} [|\rho_k - z_{k,t}|^2] \\ &= \hat{\delta}_{k,t} - 2\mathbb{E} [(\mu_{k,t} - \mu_k^*)(\rho_k - G_k(\boldsymbol{\mu}_t))] + \epsilon_k \mathbb{E} [|\rho_k - z_{k,t}|^2], \end{aligned} \tag{A-20}$$

where (a) holds since remaining budgets monotonically decrease with t , and by conditioning on the multipliers $\boldsymbol{\mu}_t$ and using the independence of $v_{k,t}$ from the multipliers $\boldsymbol{\mu}_t$ to obtain that $\mathbb{E}[Z_{k,t} | \boldsymbol{\mu}_t] = G_k(\boldsymbol{\mu}_t)$. For the second term in (A-20) one has:

$$\begin{aligned}
(\mu_{k,t} - \mu_k^*)(\rho_k - G_k(\boldsymbol{\mu}_t)) &= (\mu_{k,t} - \mu_k^*)(\rho_k - G_k(\boldsymbol{\mu}^*) + G_k(\boldsymbol{\mu}^*) - G_k(\boldsymbol{\mu}_t)) \\
&\geq (\mu_{k,t} - \mu_k^*)(G_k(\boldsymbol{\mu}^*) - G_k(\boldsymbol{\mu}_t)) ,
\end{aligned}$$

where the inequality follows because $\mu_{k,t} \geq 0$ and $\rho_k - G_k(\boldsymbol{\mu}^*) \geq 0$ and $\mu_k^*(\rho_k - G_k(\boldsymbol{\mu}^*)) = 0$. Summing over the different advertisers and Assumption 4.1, part (1) we obtain

$$\sum_{k=1}^K (\mu_{k,t} - \mu_k^*)(\rho_k - G_k(\boldsymbol{\mu}_t)) \geq \sum_{k=1}^K (\mu_{k,t} - \mu_k^*)(G_k(\boldsymbol{\mu}^*) - G_k(\boldsymbol{\mu}_t)) \geq \lambda \|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2.$$

Let $\bar{v} = \max_k \bar{v}_k$. In addition, for the third term in (A-20) we have $\mathbb{E} \left[|\rho_k - z_{k,t}|^2 \right] \leq \bar{v}_k^2 \leq \bar{v}^2$, since $\bar{v}_k \geq \rho_k \geq 0$ and $z_{k,t} \geq 0$, and since the payment is at most the bid $v_{k,t}/(1 + \mu_{k,t}) \leq v_{k,t} \leq \bar{v}_k$ because $\mu_{k,t} \geq 0$. Denoting $\bar{\epsilon} = \max_{k \in \{1, \dots, K\}} \epsilon_k$, $\underline{\epsilon} = \min_{k \in \{1, \dots, K\}} \epsilon_k$, we conclude by summing (A-20) over k that:

$$\hat{\delta}_{t+1} \leq \hat{\delta}_t - 2\lambda\delta_t + \bar{\epsilon}K\bar{v}^2 \stackrel{(a)}{\leq} (1 - 2\lambda\underline{\epsilon})\hat{\delta}_t + K\bar{v}^2\bar{\epsilon},$$

where (a) follows from $\delta_t = \sum_{k=1}^K \delta_{k,t} = \sum_{k=1}^K \epsilon_k \hat{\delta}_{k,t} \geq \underline{\epsilon} \hat{\delta}_t$ because $\delta_{k,t} \geq 0$. Lemma C.4 with $a = 2\lambda\underline{\epsilon} \leq 1$ and $b = \bar{\epsilon}K\bar{v}^2$ implies that

$$\hat{\delta}_t \leq \hat{\delta}_1 (1 - 2\lambda\underline{\epsilon})^{t-1} + \frac{K\bar{v}^2\bar{\epsilon}}{2\lambda\underline{\epsilon}}.$$

Using that $\delta_t \leq \bar{\epsilon} \hat{\delta}_t$ together with $\hat{\delta}_1 \leq \delta_1/\underline{\epsilon} \leq K\bar{\mu}^2/\underline{\epsilon}$ because $\mu_{k,t}, \mu_k^* \in [0, \bar{\mu}_k]$ and $\bar{\mu} = \max_k \bar{\mu}_k$ we obtain that

$$\delta_t \leq K\bar{\mu}^2 \frac{\bar{\epsilon}}{\underline{\epsilon}} (1 - 2\lambda\underline{\epsilon})^{t-1} + \frac{K\bar{v}^2\bar{\epsilon}^2}{2\lambda\underline{\epsilon}}. \tag{A-21}$$

Step 2. Let $\tilde{\tau}_k$ be the first auction in which the remaining budget of advertiser k is less than \bar{v}_k (at the beginning of the auction), that is, $\tilde{\tau}_k = \inf\{t \geq 1 : \tilde{B}_{k,t} < \bar{v}_k\}$. Let $\tau_k := \tilde{\tau}_k - 1$ be the last period in which the remaining budget of advertiser k is greater than \bar{v}_k . Let $\tau = \min_{k=1, \dots, K} \{\tau_k\}$. Since $v/(1 + \mu) \leq \bar{v}_k$ for any $v \in [0, \bar{v}_k]$ and $\mu \geq 0$, for any period $t \leq \tau$ the bids of all advertisers are guaranteed to be $b_{k,t}^\lambda = v_{k,t}/(1 + \mu_{k,t})$. Inequality (A-4) implies that for each bidder k , the stopping time satisfies:

$$T - \tau_k \leq \frac{\bar{\mu}_k}{\epsilon_k \rho_k} + \frac{\bar{v}_k}{\rho_k},$$

and therefore, denoting $\underline{\rho} = \min_k \rho_k$, $\bar{v} = \max_k \bar{v}_k$, and $\bar{\mu} = \max_k \bar{\mu}_k$, one has:

$$T - \tau = T - \min_{k=1,\dots,K} \{\tau_k\} = \max_{k=1,\dots,K} \{T - \tau_k\} \leq \frac{\bar{\mu}}{\underline{\epsilon}\rho} + \frac{\bar{v}}{\underline{\rho}}. \quad (\text{A-22})$$

Step 3. Putting everything together we obtain that

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}_{\mathbf{v}} [\|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2] &\stackrel{(a)}{=} \sum_{t=1}^T \mathbb{E}_{\mathbf{v}} [\|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2 \mathbf{1}\{t \leq \tau\} + \|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2 \mathbf{1}\{t > \tau\}] \\ &\stackrel{(b)}{\leq} \sum_{t=1}^T \delta_t + K\bar{\mu}^2 \mathbb{E}_{\mathbf{v}} [(T - \tau)^+] \\ &\stackrel{(c)}{\leq} K\bar{\mu}^2 \frac{\bar{\epsilon}}{\underline{\epsilon}} \sum_{t=1}^T (1 - 2\lambda\underline{\epsilon})^{t-1} + T \frac{K\bar{v}^2 \bar{\epsilon}^2}{2\lambda \underline{\epsilon}} + K\bar{\mu}^2 \left(\frac{\bar{\mu}}{\underline{\epsilon}\rho} + \frac{\bar{v}}{\underline{\rho}} \right) \\ &\stackrel{(d)}{\leq} \frac{K\bar{\mu}^2 \bar{\epsilon}}{2\lambda \underline{\epsilon}^2} + T \frac{K\bar{v}^2 \bar{\epsilon}^2}{2\lambda \underline{\epsilon}} + K\bar{\mu}^2 \left(\frac{\bar{\mu}}{\underline{\epsilon}\rho} + \frac{\bar{v}}{\underline{\rho}} \right), \end{aligned}$$

where (a) follows from conditioning on the stopping time, (b) follows from the definition of δ_t and using that $\mu_t \in [0, \bar{\mu}_k]$, (c) follows from (A-21) and (A-22), and (d) follows from $\sum_{t=1}^T (1 - 2\lambda\underline{\epsilon})^{t-1} \leq \sum_{t=0}^{\infty} (1 - 2\lambda\underline{\epsilon})^t = \frac{1}{2\lambda\underline{\epsilon}}$ because $2\lambda\underline{\epsilon} < 1$. Diving by T , the result then follows because, as $T \rightarrow \infty$, the first and third term go to zero by Assumption 4.2, part (iii), using $\underline{\epsilon} \leq \bar{\epsilon}$; and the second term goes to zero by Assumption 4.2, part (ii). \square

A.5 Proof of Theorem 4.4

In the proof we consider an alternate framework in which advertisers are also allowed to bid after depleting their budget. The main idea of the proof lies in analyzing the performance of a given advertiser, showing that its performance in the original framework (throughout its entire campaign) is close to the one it achieves in the alternate one before some advertiser runs out of budget.

Preliminaries and auxiliary results. Consider the sequence $\{(z_{k,t}, u_{k,t})\}_{t \geq 1}$ of realized expenditures and utilities of advertiser k in the alternate framework. Then, with the competing bid given by $d_{k,t} = \max_{i \neq k} v_{i,t} / (1 + \mu_{i,t})$, one has under the second-price allocation rule $z_{k,t} = \mathbf{1}\{(1 + \mu_{k,t})d_{k,t} \leq v_{k,t}\}d_{k,t}$, and $u_{k,t} = \mathbf{1}\{(1 + \mu_{k,t})d_{k,t} \leq v_{k,t}\}(v_{k,t} - d_{k,t})$. Let $\tilde{B}_{k,t} = B_k - \sum_{s=1}^{t-1} z_{s,k}$ denote the evolution of the k^{th} advertiser's budget at the beginning of period t in the alternate framework.

Let $\tau = \inf\{t \geq 1 : \tilde{B}_{k,t+1} < \bar{v}_k \text{ for some } k = 1, \dots, K\}$ be the last period in which the remaining budget of all advertisers is larger than \bar{v}_k . Since $v/(1 + \mu) \leq \bar{v}_k$ for any $v \in [0, \bar{v}_k]$ and $\mu \geq 0$, for any period $t \leq \tau$ the bids of all advertisers are guaranteed to be $b_{k,t}^A = v_{k,t}/(1 + \mu_{k,t})$. Denoting $\underline{\rho} = \min_k \rho_k$, $\bar{v} = \max_k \bar{v}_k$, and $\bar{\mu} = \max_k \bar{\mu}_k$, inequality (A-22) implies that

$$T - \tau \leq \frac{\bar{\mu}}{\underline{\epsilon}\rho} + \frac{\bar{v}}{\rho}.$$

A key step in the proof involves showing that the utility-per-auction collected by the advertiser is “close” relative to $\Psi_k(\mu^*)$. The next result bounds the performance gap of the adaptive pacing strategy relative to $\Psi_k(\mu^*)$, in terms of the expected squared error of multiplier. Denote by $\boldsymbol{\mu}_t \in \mathbb{R}^K$ the (random) vector such that $\mu_{i,t}$ is the multiplier used by advertiser i at time t .

Lemma A.3. *Let each advertiser k follow the adaptive pacing strategy A in the alternate framework where budgets are not enforced. Then, there exists a constant $C_3 > 0$ such that:*

$$\mathbb{E}[u_{k,t}] \geq \frac{\Psi_k(\boldsymbol{\mu}^*)}{T} - C_3 (\mathbb{E}[\|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2])^{1/2}.$$

We note that $C_3 = (2\bar{v} + \bar{\mu}\bar{v}^2\bar{f})K^{1/2}$. The proof of Lemma A.3 is established by writing the utility in terms of the dual function and the complementary slackness condition, and then using that the dual and functions are Lipschitz continuous as argued in Lemma C.2. The proof of this result is deferred to Appendix B.

Proving the result. Let each advertiser k follow the adaptive pacing strategy with step size ϵ_k . The performance of both the original and the alternate systems coincide until time τ , and therefore:

$$\Pi_k^A \stackrel{(a)}{\geq} \mathbb{E} \left[\sum_{t=1}^{\tau \wedge T} u_{k,t} \right] \stackrel{(b)}{\geq} \mathbb{E} \left[\sum_{t=1}^T u_{k,t} \right] - \bar{v}_k \mathbb{E}[(T - \tau)^+],$$

where (a) follows from discarding all auctions after the time some advertiser runs out of budget; and (b) follows from $0 \leq u_{k,t} \leq \bar{v}_k$. Summing the lower bound on the expected utility-per-auction in Lemma A.3 over $t = 1, \dots, T$ one has:

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T u_{k,t} \right] &\geq \Psi_k(\boldsymbol{\mu}^*) - C_3 \sum_{t=1}^T (\mathbb{E}[\|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2])^{1/2} \\ &\stackrel{(a)}{\geq} \Psi_k(\boldsymbol{\mu}^*) - C_3 C_1^{1/2} \left(\frac{\bar{\epsilon}}{\underline{\epsilon}} \right)^{1/2} \sum_{t=1}^T (1 - 2\lambda\underline{\epsilon})^{(t-1)/2} - C_3 C_2^{1/2} \frac{\bar{\epsilon}}{\underline{\epsilon}^{1/2}} T \end{aligned} \quad (\text{A-23})$$

$$\stackrel{(b)}{\geq} \Psi_k(\boldsymbol{\mu}^*) - \frac{C_3 C_1^{1/2}}{\lambda} \frac{\bar{\epsilon}^{1/2}}{\underline{\epsilon}^{3/2}} - C_3 C_2^{1/2} \frac{\bar{\epsilon}}{\underline{\epsilon}^{1/2}} T, \quad (\text{A-24})$$

where (a) follows from Theorem 4.3 and $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$ for $x, y \geq 0$, and (b) follows from $\sum_{t=1}^T (1 - 2\lambda\underline{\epsilon})^{(t-1)/2} \leq \sum_{t=0}^{\infty} (1 - 2\lambda\underline{\epsilon})^{t/2} = \frac{1}{1 - (1 - 2\lambda\underline{\epsilon})^{1/2}} \leq \frac{1}{\lambda\underline{\epsilon}}$ because $1 - (1 - x)^{1/2} \geq x/2$ for $x \in [0, 1]$. We use the bound in (A-22) to bound the truncated expectation as follows:

$$\mathbb{E} [(T - \tau)^+] \leq \frac{\bar{v}\bar{\mu}}{\underline{\epsilon}\underline{\rho}} + \frac{\bar{v}^2}{\underline{\rho}}. \quad (\text{A-25})$$

Combining (A-23) and (A-25) we obtain that

$$\begin{aligned} \Pi_k^A &\geq \mathbb{E} \left[\sum_{t=1}^T u_{k,t} \right] - \bar{v}_k \mathbb{E} [(T - \tau)^+] \\ &\geq \Psi_k(\boldsymbol{\mu}^*) - \frac{C_3 C_1^{1/2} \bar{\epsilon}^{1/2}}{\lambda \underline{\epsilon}^{3/2}} - C_3 C_2^{1/2} \frac{\bar{\epsilon}}{\underline{\epsilon}^{1/2}} T - \frac{\bar{\mu}}{\underline{\epsilon}\underline{\rho}} - \frac{\bar{v}}{\underline{\rho}}. \end{aligned} \quad (\text{A-26})$$

The dependence of the payoff gap on the number of advertisers K is of the same as the dependencies of the products $C_3 C_1^{1/2}$ and $C_3 C_2^{1/2}$, which are of order K . This concludes the proof. \square

A.6 Proof of Theorem 5.2

As in the proof of Theorem 4.4, we consider an alternate framework in which advertisers are allowed to bid even after budget depletion, without any utility gained. The performance of a deviating advertiser (indexed k) in the original framework is equal to the one in the alternate framework up to the first time some advertiser runs out of budget. As an adaptive pacing strategy does not run out of budget too early, the main idea of the proof lies in analyzing the performance of the deviating advertiser in the alternate framework. We first argue that when the number of competitors is large, the expected squared error of the competitors' multipliers relative to the vector $\boldsymbol{\mu}^*$ is small, since the impact of the deviating advertiser on its competitors is limited. We then show that the benefit of deviating to any other strategy is small when competitors' multipliers are "close" to $\boldsymbol{\mu}^*$.

Preliminaries and auxiliary results. Consider the sequence $\{(z_{k,t}, u_{k,t})\}_{t \geq 1}$ of realized expenditures and utilities of advertiser k in the alternate framework, and let $\{b_{k,t}^\beta\}_{t \geq 1}$ be the bids of advertiser k . Then, with the competing bid given by $d_{k,t} = \max_{i \neq k, m_{k,t} = m_{i,t}} v_{i,t} / (1 + \mu_{i,t})$, one has under the second-price allocation rule that $z_{k,t} = \mathbf{1}\{d_{k,t} \leq b_{k,t}^\beta\} d_{k,t}$, and $u_{k,t} = \mathbf{1}\{d_{k,t} \leq b_{k,t}^\beta\} (v_{k,t} - d_{k,t})$. The competing bid faced by an advertiser $i \neq k$ is given by $d_{i,t} = b_{k,t}^\beta \vee \max_{j \neq k, i; m_{k,t} = m_{j,t}} v_{j,t} / (1 + \mu_{j,t})$. The history available at time t to advertiser k in the model described in §5 is defined by

$$\mathcal{H}_{k,t} := \sigma \left(\left\{ m_{k,\tau}, v_{k,\tau}, b_{k,\tau}, z_{k,\tau}, u_{k,\tau} \right\}_{\tau=1}^{t-1}, m_{k,t}, v_{k,t}, y \right)$$

for any $t \geq 2$, with $\mathcal{H}_{k,1} := \sigma(m_{k,1}, v_{k,1}, y)$. In what follows we provide few auxiliary results that we use in the proof. The complete proofs of these results are deferred to Appendix B.

Let τ_i be the first auction in which the remaining budget of advertiser $i \neq k$ is less than \bar{v}_i (just before the auction); that is, $\tau_i = \inf\{t \geq 1 : \tilde{B}_{i,t} < \bar{v}_i\}$, and let $\tau = \min_{i \neq k} \{\tau_i\}$. Since for each advertiser $i \neq k$ the bid satisfies $v/(1 + \mu) \leq \bar{v}_i$ for any $v \in [0, \bar{v}_i]$ and $\mu \geq 0$, in each period $t \leq \tau$ the bid is $b_{i,t} = v_{i,t}/(1 + \mu_{i,t})$. Inequality (A-22) implies that the stopping time satisfies

$$T - \tau \leq \frac{\bar{\mu}}{\underline{\epsilon}\rho} + \frac{\bar{v}}{\rho}, \quad (\text{A-27})$$

with $\underline{\rho} = \min_{i \neq k} \rho_i$, $\bar{v} = \max_{i \neq k} \bar{v}_i$, and $\bar{\mu} = \max_{i \neq k} \bar{\mu}_i$. First, we show that the mean squared errors of the estimated multipliers at period t can be bounded in terms of the minimum and maximum step sizes, the number of players and $\mathbf{a}_k = (a_{k,i})_{i \neq k} \in [0, 1]^{K-1}$. Denote by $\boldsymbol{\mu}_t \in \mathbb{R}^K$ the (random) vector such that $\mu_{k,t} = \mu_k^*$ and $\mu_{i,t}$ is the multiplier used by advertiser $i \neq k$ at time t .

Lemma A.4. *Suppose that Assumption 4.1 holds and let $\boldsymbol{\mu}^*$ be the profile of multipliers defined by (5). Let each advertiser $i \neq k$ follow the adaptive pacing strategy A with $\bar{\epsilon} \leq 1/(2\lambda)$ and suppose that advertiser k uses some strategy $\beta \in \mathcal{B}_k^{\text{CI}}$. Then, there exist positive constants C_1, C_2, C_3 independent of T and K , such that for any $t \in \{1, \dots, T\}$:*

$$\mathbb{E} [\|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2] \leq C_1 K \frac{\bar{\epsilon}}{\underline{\epsilon}} (1 - \lambda \underline{\epsilon})^{t-1} + C_2 K \frac{\bar{\epsilon}^2}{\underline{\epsilon}} + C_3 \|\mathbf{a}_k\|_2^2 \frac{\bar{\epsilon}^2}{\underline{\epsilon}^2}.$$

Lemma A.4 implies that the vector of multipliers $\boldsymbol{\mu}_t$ selected by advertisers converges in L_2 to the vector $\boldsymbol{\mu}^*$ under some assumptions on the step sizes and how often advertiser k interacts with its competitors. The next result bounds the Lagrangian utility per auction of an adaptive pacing strategy relative to $\Psi_k(\boldsymbol{\mu}^*)$, in terms of the multipliers' expected squared error.

Lemma A.5. *Suppose that Assumption 4.1 holds and let $\boldsymbol{\mu}^*$ be the profile of multipliers defined by (5). Let each advertiser $i \neq k$ follow the adaptive pacing strategy A. Then, there exists a constant $C_4 > 0$ such that for any $t \in \{1, \dots, T\}$:*

$$\mathbb{E} [u_{k,t} - \mu_k^* z_{k,t} + \rho_k] \leq \frac{\Psi_k(\boldsymbol{\mu}^*)}{T} + C_4 \|\mathbf{a}_k\|_2 \cdot \left(\mathbb{E} \|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2 \right)^{1/2}.$$

Proving the result. Let $\boldsymbol{\mu}^*$ be the profile of multipliers defined by (5). Let each advertiser $i \neq k$ follow the adaptive pacing strategy with step size ϵ_k , except advertiser k who implements a strategy $\beta \in \mathcal{B}_k^{\text{CI}}$. Both the original and the alternate systems coincide up to time τ , and thus:

$$\Pi_k^{\beta, \mathbf{A}-k} \stackrel{(a)}{\leq} \mathbb{E} \left[\sum_{t=1}^T u_{k,t} \right] + \bar{v}_k \mathbb{E} [(T - \tau)^+] \stackrel{(b)}{\leq} \sum_{t=1}^T \mathbb{E} [u_{k,t} - \mu_k^* z_{k,t} + \rho_k] + \bar{v}_k \mathbb{E} [(T - \tau)^+],$$

where (a) follows from adding all auctions after time τ in the alternate system, and using that the utility of each auction in the original system satisfy $0 \leq u_{k,t} \leq \bar{v}_k$; and (b) follows from adding the constraint $\sum_{t=1}^T z_{k,t} \leq B_k$ to the objective with a Lagrange multiplier μ_k^* , because the strategy β is budget-feasible in the alternate framework. Summing the lower bound on the expected utility-per-auction in Lemma A.5 over $t = 1, \dots, T$ one has:

$$\sum_{t=1}^T \mathbb{E} [u_{k,t} - \mu_k^* z_{k,t} + \rho_k] \leq \Psi_k(\boldsymbol{\mu}^*) + C_4 \|\mathbf{a}_k\|_2 \sum_{t=1}^T \left(\mathbb{E} \|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2 \right)^{1/2}.$$

The sum in the second term can be upper bounded by

$$\begin{aligned} \sum_{t=1}^T \left(\mathbb{E} \|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2 \right)^{1/2} &\stackrel{(a)}{\leq} C_1^{1/2} K^{1/2} \left(\frac{\bar{\epsilon}}{\underline{\epsilon}} \right)^{1/2} \sum_{t=1}^T (1 - \lambda \underline{\epsilon})^{(t-1)/2} + C_2^{1/2} K^{1/2} \frac{\bar{\epsilon}}{\underline{\epsilon}^{1/2}} T + C_3^{1/2} \|\mathbf{a}_k\|_2 \frac{\bar{\epsilon}}{\underline{\epsilon}} T \\ &\stackrel{(b)}{\leq} \frac{2C_1^{1/2}}{\lambda} K^{1/2} \frac{\bar{\epsilon}^{1/2}}{\underline{\epsilon}^{3/2}} + C_2^{1/2} K^{1/2} \frac{\bar{\epsilon}}{\underline{\epsilon}^{1/2}} T + C_3^{1/2} \|\mathbf{a}_k\|_2 \frac{\bar{\epsilon}}{\underline{\epsilon}} T, \end{aligned}$$

where (a) follows from Lemma A.4 and $\sqrt{x+y+z} \leq \sqrt{x} + \sqrt{y} + \sqrt{z}$ for $x, y, z \geq 0$, and (b) follows from $\sum_{t=1}^T (1 - \lambda \underline{\epsilon})^{(t-1)/2} \leq \sum_{t=0}^{\infty} (1 - \lambda \underline{\epsilon})^{t/2} = \frac{1}{1 - (1 - \lambda \underline{\epsilon})^{1/2}} \leq \frac{2}{\lambda \underline{\epsilon}}$ because $1 - (1 - x)^{1/2} \geq x/2$ for $x \in [0, 1]$. Therefore, we obtain that

$$\begin{aligned} \sum_{t=1}^T \mathbb{E} [u_{k,t} - \mu_k^* z_{k,t} + \rho_k] &\leq \Psi_k(\boldsymbol{\mu}^*) + \frac{2C_4 C_1^{1/2}}{\lambda} \|\mathbf{a}_k\|_2 K^{1/2} \frac{\bar{\epsilon}^{1/2}}{\underline{\epsilon}^{3/2}} \\ &\quad + C_4 C_2^{1/2} \|\mathbf{a}_k\|_2 K^{1/2} \frac{\bar{\epsilon}}{\underline{\epsilon}^{1/2}} T + C_4 C_3^{1/2} \|\mathbf{a}_k\|_2^2 \frac{\bar{\epsilon}}{\underline{\epsilon}} T. \end{aligned} \quad (\text{A-28})$$

We use the bound in (A-27) to bound the truncated expectation as follows:

$$\mathbb{E} [(T - \tau)^+] \leq \frac{\bar{\mu}}{\underline{\epsilon} \rho} + \frac{\bar{v}}{\underline{\rho}}. \quad (\text{A-29})$$

Combining (A-28) and (A-29) one obtains that

$$\begin{aligned} \Pi_k^{\beta, \mathbf{A}-k} &\leq \mathbb{E} \left[\sum_{t=1}^T u_{k,t} \right] + \bar{v}_k \mathbb{E} [(T - \tau)^+] \\ &\leq \Psi_k(\boldsymbol{\mu}^*) + \frac{2C_4 C_1^{1/2}}{\lambda} \|\mathbf{a}_k\|_2 K^{1/2} \frac{\bar{\epsilon}^{1/2}}{\underline{\epsilon}^{3/2}} \\ &\quad + C_4 C_2^{1/2} \|\mathbf{a}_k\|_2 K^{1/2} \frac{\bar{\epsilon}}{\underline{\epsilon}^{1/2}} T + C_4 C_3^{1/2} \|\mathbf{a}_k\|_2^2 \frac{\bar{\epsilon}}{\underline{\epsilon}} T + \frac{\bar{\mu} \bar{v}}{\underline{\epsilon} \rho} + \frac{\bar{v}^2}{\underline{\rho}}. \end{aligned}$$

The result follows from combining the last inequality with (A-26). \square

B Proofs of key lemmas

Proof of Lemma A.1. To simplify notation we drop the dependence on k . Let $\mathbb{P}\{\mathbf{d}^i\}$ be the probability that competing bid sequence \mathbf{d}^i is chosen, for $i = 1, \dots, m$, and let β be a (potentially randomized) strategy. We have:

$$\begin{aligned} \sup_{\substack{\mathbf{v} \in [0, \bar{v}]^T \\ \mathbf{d} \in \mathbb{R}_+^T}} \mathbb{E}^\beta \left[R_\gamma^\beta(\mathbf{v}; \mathbf{d}) \right] &\stackrel{(a)}{=} \sum_{i=1}^m \mathbb{P}\{\mathbf{d}^i\} \sup_{\substack{\mathbf{v} \in [0, \bar{v}]^T \\ \mathbf{d} \in \mathbb{R}_+^T}} \mathbb{E}^\beta \left[R_\gamma^\beta(\mathbf{v}; \mathbf{d}) \right] \stackrel{(b)}{\geq} \sum_{i=1}^m \mathbb{P}\{\mathbf{d}^i\} \mathbb{E}^\beta \left[R_\gamma^\beta(\mathbf{v}'; \mathbf{d}^i) \right] \\ &\stackrel{(c)}{=} \mathbb{E}^\beta \left[\sum_{i=1}^m \mathbb{P}\{\mathbf{d}^i\} R_\gamma^\beta(\mathbf{v}'; \mathbf{d}^i) \right] \stackrel{(d)}{\geq} \inf_{\beta \in \tilde{\mathcal{B}}} \sum_{i=1}^m \mathbb{P}\{\mathbf{d}^i\} R_\gamma^\beta(\mathbf{v}'; \mathbf{d}^i), \end{aligned}$$

where: (a) follows from $\sum_i \mathbb{P}\{\mathbf{d}^i\} = 1$; (b) holds since \mathbf{v}' and \mathbf{d}^i are feasible realizations; (c) follows from Fubini's Theorem since one has $|R_\gamma^\beta(\mathbf{v}; \mathbf{d})| \leq \frac{1}{T} |\pi^H(\mathbf{v}; \mathbf{d})| + \frac{\gamma}{T} |\pi^\beta(\mathbf{v}; \mathbf{d})| \leq \bar{v}(1 + \gamma)$ because no strategy (even in hindsight) can achieve more than $T\bar{v}$; and (d) follows because any randomized strategy can be thought of as a probability distribution over deterministic algorithms. \square

Proof of Lemma A.2. To simplify notation we drop the dependence on k . We use the worst-case instance structure detailed in the proof of Theorem 3.1. Fix any deterministic bidding strategy $\beta \in \tilde{\mathcal{B}}$. Since β is deterministic one has that $\pi^\beta(\mathbf{d}, \mathbf{v}) = \sum_{t=1}^T \mathbf{1}\{d_t \leq b_t^\beta\}(v_t - d_t)$ for any vectors \mathbf{d} and \mathbf{v} where b_t^β is the bid dictated by β at time t .

Let $\bar{\mathbf{v}} = (\bar{v}, \dots, \bar{v})$ be the valuation sequence and $\mathbf{d}^i \in \mathcal{D}$ be the sequence of competing bids. We denote by b_t^i the bid in period t under this input and β . We also denote the corresponding expenditure by $z_t^i := \mathbf{1}\{d_t^i \leq b_t^i\}d_t^i$, and the corresponding net utility by $u_t^i := \mathbf{1}\{d_t^i \leq b_t^i\}(\bar{v} - d_t^i)$. We further denote the history of decisions and observations under $\bar{\mathbf{v}}$, \mathbf{d}^i and β by $\mathcal{H}_t^i := \sigma\left(\left\{v_\tau^i, b_\tau^i, z_\tau^i, u_\tau^i\right\}_{\tau=1}^{t-1}, \bar{v}\right)$ for any $t \geq 2$, with $\mathcal{H}_1^i := \sigma(\bar{v})$.

We now define the sequence \mathbf{x} . For each $j \in \{1, \dots, m\}$, define:

$$x_j := \sum_{t=(j-1)\lfloor \frac{T}{m} \rfloor + 1}^{j\lfloor \frac{T}{m} \rfloor} \mathbf{1}\{d_t^1 \leq b_t^1\},$$

where we denote by b_t^1 the bid at time t under history \mathcal{H}_t^1 . Then, x_j is the number of auctions won by β throughout the j 'th batch of $\lfloor T/m \rfloor$ auctions when the vector of best competitors' bids

is \mathbf{d}^1 and the vector of valuations is $\bar{\mathbf{v}}$. Since \mathbf{d}^1 , $\bar{\mathbf{v}}$ and β are deterministic, so is \mathbf{x} , and clearly $\mathbf{x} \in \mathcal{X}$ because β is budget-feasible. Denote by β^x a strategy that for each $j \in \{1, \dots, m\}$ bids d_j in the first x_j auctions of batch j . When competing bids are d_j in the j 'th batch, this strategy is guaranteed to win x_j auctions.

We next show that β^x achieves the same performance as β . First, assume that the vector of competing bids is \mathbf{d}^1 . Then, by construction:

$$\pi^\beta(\bar{\mathbf{v}}, \mathbf{d}^1) = \sum_{t=1}^T \mathbf{1}\{d_t^1 \leq b_t^1\}(\bar{v} - d_t^1) = \sum_{j=1}^m x_j(\bar{v} - d_j) = \pi^{\beta^x}(\bar{\mathbf{v}}, \mathbf{d}^1),$$

where d_j denotes the competing bid of the j 'th batch of the competing bid sequence \mathbf{d}^1 . For any $\mathbf{d}^i \in \mathcal{D} \setminus \{\mathbf{d}^1\}$, define $\tau^i := \inf \{t \geq 1 : d_t^i \neq d_t^1\}$ to be the first time that the competing bid sequence is different than \mathbf{d}^1 . Let $m^i \in \{1, \dots, m\}$ be the number of batches that the competing bid sequence \mathbf{d}^i has in common with \mathbf{d}^1 . Since each batch has $\lfloor T/m \rfloor$ items, we have that $\tau^i = m^i \lfloor T/m \rfloor + 1$. The sequences \mathbf{d}^1 and \mathbf{d}^i are identical, thus indistinguishable, up to time τ^i . Therefore, the bids of any deterministic strategy coincide up to time τ^i under histories \mathcal{H}_t^1 and \mathcal{H}_t^i . Then, one has:

$$\pi^\beta(\bar{\mathbf{v}}, \mathbf{d}^i) \stackrel{(a)}{=} \sum_{t=1}^{\tau^i-1} \mathbf{1}\{d_t^i \leq b_t^i\}(\bar{v} - d_t^i) \stackrel{(b)}{=} \sum_{t=1}^{\tau^i-1} \mathbf{1}\{d_t^1 \leq b_t^1\}(\bar{v} - d_t^1) \stackrel{(c)}{=} \sum_{j=1}^{m^i} x_j(\bar{v} - d_j) = \pi^{\beta^x}(\bar{\mathbf{v}}, \mathbf{d}^i),$$

where (a) follows because all items in periods $t \in \{\tau^i, \dots, T\}$ have zero utility under \mathbf{d}^i , (b) follows because the competing bid sequences and bids are equal during periods $t \in \{1, \dots, \tau^i - 1\}$, (c) follows from our definition of \mathbf{x} and using that sequence \mathbf{d}^i has m^i batches with nonzero utility. We have thus established that $\pi^{\beta^x}(\bar{\mathbf{v}}, \mathbf{d}^i) = \pi^\beta(\bar{\mathbf{v}}, \mathbf{d}^i)$ for any $\mathbf{d} \in \mathcal{D}$. This concludes the proof. \square

Proof of Lemma A.3. Let each advertiser k follow an adaptive pacing strategy, and define $\delta_t = \sum_{k=1}^K \mathbb{E}_{\mathbf{v}} [|\mu_{k,t} - \mu_k^*|^2]$. Denote $\bar{\Psi}_k(\boldsymbol{\mu}) = \Psi_k(\boldsymbol{\mu})/T$, and let $\hat{\Psi}_{k,t}(\boldsymbol{\mu}) = (v_{k,t} - (1 + \mu_{k,t})d_{k,t})^+ + \mu_{k,t}\rho_k$ be normalized empirical dual objective function for advertiser k , with $d_{k,t} = \max_{i \neq k} v_{i,t}/(1 + \mu_{i,t})$. By the structure of the strategy, the utility of advertiser k from the t^{th} auction can be written as:

$$\begin{aligned} u_{k,t} &= \mathbf{1}\{d_{k,t}(1 + \mu_{k,t}) \leq v_{k,t}\}(v_{k,t} - d_{k,t}) \\ &= (v_{k,t} - (1 + \mu_{k,t})d_{k,t})^+ + \mu_{k,t}z_{k,t} = \hat{\Psi}_{k,t}(\boldsymbol{\mu}_t) + \mu_{k,t}(z_{k,t} - \rho_k). \end{aligned}$$

Taking expectations we obtain:

$$\mathbb{E}[u_{k,t}] \stackrel{(a)}{=} \mathbb{E} \left[\mathbb{E} \left[\hat{\Psi}_{k,t}(\boldsymbol{\mu}_t) + \mu_{k,t} (z_{k,t} - \rho_k) \mid \boldsymbol{\mu}_t \right] \right] \stackrel{(b)}{=} \mathbb{E} [\bar{\Psi}_k(\boldsymbol{\mu}_t)] + \mathbb{E} [\mu_{k,t} (G_k(\boldsymbol{\mu}_t) - \rho_k)],$$

where (a) follows from the linearity of expectation and conditioning on the multipliers $\boldsymbol{\mu}_t$, and (b) holds since $\{v_{k,t}\}_{k=1}^K$ are independent of the multipliers $\boldsymbol{\mu}_t$. Let $\bar{v} = \max_k \bar{v}_k$. Since the dual objective is Lipschitz continuous (Lemma C.2, part (ii)) one has:

$$\bar{\Psi}_k(\boldsymbol{\mu}_t) \geq \bar{\Psi}_k(\boldsymbol{\mu}^*) - \bar{v} \|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_1,$$

for all $t = 1, \dots, T$ because $a_{k,i} \in [0, 1]$. Since the expenditure is Lipschitz continuous (Lemma C.2, part (iii)) one has:

$$\begin{aligned} \mu_{k,t} (G_k(\boldsymbol{\mu}_t) - \rho_k) &= \mu_{k,t} (G_k(\boldsymbol{\mu}_t) - G_k(\boldsymbol{\mu}^*)) + (\mu_{k,t} - \mu_k^*) (G_k(\boldsymbol{\mu}^*) - \rho_k) + \mu_k^* (G_k(\boldsymbol{\mu}^*) - \rho_k) \\ &\stackrel{(a)}{\geq} -(\bar{\mu}_k \bar{v}^2 \bar{f} + \bar{v}) \|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_1, \end{aligned}$$

for all $t = 1, \dots, T$, where (a) follows from $\mu_{k,t} \leq \bar{\mu}_k$, $|G_k(\boldsymbol{\mu}^*) - \rho_k| \leq \bar{v}$ together with $\mu_k^* (G_k(\boldsymbol{\mu}^*) - \rho_k) = 0$, which follows the characterization of $\boldsymbol{\mu}^*$ in (5). In addition, one has:

$$\sum_{k=1}^K \mathbb{E} |\mu_{k,t} - \mu_k^*| \stackrel{(a)}{\leq} \mathbb{E} \left[\left(K \sum_{k=1}^K |\mu_{k,t} - \mu_k^*|^2 \right)^{1/2} \right] \stackrel{(b)}{\leq} K^{1/2} \left(\sum_{k=1}^K \mathbb{E} [|\mu_{k,t} - \mu_k^*|^2] \right)^{1/2} = K^{1/2} \delta_t^{1/2}$$

where (a) follows from $\sum_{i=1}^n |y_i| \leq (n \sum_{i=1}^n y_i^2)^{1/2}$, and (b) follows from Jensen's inequality. Together, we obtain:

$$\mathbb{E}[u_{k,t}] \geq \bar{\Psi}_k(\boldsymbol{\mu}^*) - (2\bar{v} + \bar{\mu}_k \bar{v}^2 \bar{f}) K^{1/2} \delta_t^{1/2}.$$

This concludes the proof. \square

Proof of Lemma A.4. Denote by $\boldsymbol{\mu}_t \in \mathbb{R}^K$ the (random) vector such that $\mu_{k,t} = \mu_k^*$ and $\mu_{i,t}$ is the multiplier used by advertiser $i \neq k$ at time t . Fix an advertiser $i \neq k$. Define $\delta_{i,t} := \mathbb{E} [(\mu_{i,t} - \mu_i^*)^2]$ and $\delta_i := \sum_{i \neq k} \delta_{i,t}$. Similarly, define $\hat{\delta}_{i,t} := \delta_{i,t}/\epsilon_i$ and $\hat{\delta}_t := \sum_{i \neq k} \delta_{i,t}$. We obtain from (A-20) in the proof of Theorem 4.3 that:

$$\begin{aligned} \hat{\delta}_{i,t+1} &\leq \hat{\delta}_{i,t} - 2\mathbb{E} [(\mu_{i,t} - \mu_i^*) (\rho_i - z_{i,t})] + \epsilon_i \mathbb{E} [|\rho_i - z_{i,t}|^2] \\ &= \hat{\delta}_{i,t} - 2\mathbb{E} [(\mu_{i,t} - \mu_i^*) (\rho_i - \mathbb{E}[z_{i,t} \mid \boldsymbol{\mu}_t])] + \epsilon_i \mathbb{E} [|\rho_i - z_{i,t}|^2], \end{aligned} \tag{B-1}$$

where the equality follows by conditioning on $\boldsymbol{\mu}_t$. Let $\bar{v} = \max_{k \neq i} \bar{v}_i$. The third term satisfies

$\mathbb{E} \left[|\rho_i - z_{i,t}|^2 \right] \leq \bar{v}^2$, since $\rho_i, z_{i,t} \in [0, \bar{v}]$. Proceeding to bound the second term, recall that the payment of advertiser i is $z_{i,t} = \mathbf{1}\{d_{i,t} \leq v_{i,t}/(1 + \mu_{i,t})\}d_{i,t}$ where the competing bid faced by the advertiser is given by

$$d_{i,t} = \begin{cases} b_{k,t}^\beta \vee d_{i \setminus k,t}, & \text{if } m_{k,t} = m_{i,t}, \\ d_{i \setminus k,t}, & \text{if } m_{k,t} \neq m_{i,t}, \end{cases}$$

where $d_{i \setminus k,t} = \max_{j \neq k, i; m_{j,t} = m_{i,t}} \{v_{j,t}/(1 + \mu_{j,t})\}$ denotes the maximum competing bid faced by advertiser i when advertiser k is excluded. Recall that for a fixed vector $\boldsymbol{\mu} \in \mathbb{R}_+^K$, the expected expenditure function is given by $G_i(\boldsymbol{\mu}) = \mathbb{E} \left[\tilde{d}_i \mathbf{1} \left\{ \tilde{d}_i \leq v_i/(1 + \mu_i) \right\} \right]$ where $\tilde{d}_i = \max_{j \neq i; m_j = m_i} \{v_j/(1 + \mu_j)\}$. For the second term in (B-1) one has:

$$\begin{aligned} (\mu_{i,t} - \mu_i^*) (\rho_i - \mathbb{E}[z_{i,t} | \boldsymbol{\mu}_t]) &= (\mu_{i,t} - \mu_i^*) (\rho_i - G_i(\boldsymbol{\mu}^*) + G_i(\boldsymbol{\mu}^*) - G_i(\boldsymbol{\mu}_t) + G_i(\boldsymbol{\mu}_t) - \mathbb{E}[z_{i,t} | \boldsymbol{\mu}_t]) \\ &\geq (\mu_{i,t} - \mu_i^*) (G_i(\boldsymbol{\mu}^*) - G_i(\boldsymbol{\mu}_t)) - |\mu_{i,t} - \mu_i^*| \cdot |G_i(\boldsymbol{\mu}_t) - \mathbb{E}[z_{i,t} | \boldsymbol{\mu}_t]|, \end{aligned}$$

where the inequality follows because $\mu_{i,t} \geq 0$ and $\rho_i - G_i(\boldsymbol{\mu}^*) \geq 0$ and $\mu_i^* (\rho_i - G_i(\boldsymbol{\mu}^*)) = 0$, together with $xy \geq -|x| \cdot |y|$ for $x, y \in \mathbb{R}$. Because values are independent, and advertisers i and k compete only when $m_{k,t} = m_{i,t}$, we obtain that

$$|G_i(\boldsymbol{\mu}_t) - \mathbb{E}[z_{i,t} | \boldsymbol{\mu}_t]| = \left| \mathbb{E} \left[\left(\tilde{d}_{i,t} \mathbf{1} \{ \tilde{d}_{i,t} \leq b_{i,t} \} - d_{i,t} \mathbf{1} \{ d_{i,t} \leq b_{i,t} \} \right) \mathbf{1} \{ m_{k,t} = m_{i,t} \} \right] \right| \leq \bar{v} a_{k,i},$$

where the equality follows because the bid of advertiser $i \neq k$ is $b_{i,t} = v_{i,t}/(1 + \mu_{i,t})$, and $d_{i,t} = b_{k,t}^\beta \vee d_{i \setminus k,t}$ and $\tilde{d}_{i,t} = (v_{k,t}/(1 + \mu_k^*)) \vee d_{i \setminus k,t}$ when $m_{k,t} = m_{i,t}$; and the inequality follows because the expenditure is at most the bid and $b_{i,t} \leq \bar{v}$ together with $\mathbb{P}\{m_{k,t} = m_{i,t}\} = a_{k,i}$. Summing over the different advertisers and using Assumption 4.1, part (1) we obtain by

$$\begin{aligned} \sum_{i \neq k} (\mu_{i,t} - \mu_i^*) (\rho_i - \mathbb{E}[z_{i,t} | \boldsymbol{\mu}_t]) &\geq \sum_{i=1}^K (\mu_{i,t} - \mu_i^*) (G_i(\boldsymbol{\mu}^*) - G_i(\boldsymbol{\mu}_t)) - \bar{v} \sum_{i \neq k} a_{k,i} |\mu_{i,t} - \mu_i^*| \\ &\geq \lambda \|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2 - \bar{v} \|\mathbf{a}_k\|_2 \cdot \|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2, \end{aligned}$$

where the last inequality follows from Cauchy-Schwarz inequality. Denoting $\bar{\epsilon} = \max_{k \in \{1, \dots, K\}} \epsilon_k$, $\underline{\epsilon} = \min_{k \in \{1, \dots, K\}} \epsilon_k$, we conclude by summing (A-20) over k that:

$$\hat{\delta}_{t+1} \stackrel{(a)}{\leq} \hat{\delta}_t - 2\lambda \delta_t + \bar{\epsilon} K \bar{v}^2 + 2\bar{v} \|\mathbf{a}_k\|_2 \delta_t^{1/2} \stackrel{(b)}{\leq} (1 - 2\lambda \underline{\epsilon}) \hat{\delta}_t + K \bar{v}^2 \bar{\epsilon} + 2\bar{v} \|\mathbf{a}_k\|_2 \bar{\epsilon}^{1/2} \hat{\delta}_t^{1/2},$$

where (a) follows because $\mathbb{E}\|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2 \leq (\mathbb{E}\|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2)^{1/2} = \delta_t^{1/2}$ by Jensen's Inequality, and (b) follows from $\delta_t = \sum_{i \neq k} \delta_{i,t} = \sum_{i \neq k} \epsilon_i \hat{\delta}_{i,t} \geq \underline{\epsilon} \hat{\delta}_t$ because $\delta_{i,t} \geq 0$ and $\delta_t \leq \bar{\epsilon} \hat{\delta}_t$. Lemma C.5 with $a = 2\lambda\underline{\epsilon} \leq 1$, $b = \bar{\epsilon}K\bar{v}^2$ and $c = 2\bar{v}\|\mathbf{a}_k\|_2\bar{\epsilon}^{1/2}$ implies that

$$\hat{\delta}_t \leq \hat{\delta}_1 (1 - \lambda\underline{\epsilon})^{t-1} + \frac{K\bar{v}^2\bar{\epsilon}}{\lambda\underline{\epsilon}} + \left(\frac{\bar{v}\|\mathbf{a}_k\|_2\bar{\epsilon}^{1/2}}{\lambda\underline{\epsilon}} \right)^2.$$

Using that $\delta_t \leq \bar{\epsilon}\hat{\delta}_t$ together with $\hat{\delta}_1 \leq \delta_1/\underline{\epsilon} \leq K\bar{\mu}^2/\underline{\epsilon}$ because $\mu_{i,t}, \mu_i^* \in [0, \bar{\mu}_k]$ and $\bar{\mu} = \max_k \bar{\mu}_k$ we obtain that

$$\delta_t \leq K\bar{\mu}^2 \frac{\bar{\epsilon}}{\underline{\epsilon}} (1 - \lambda\underline{\epsilon})^{t-1} + \frac{K\bar{v}^2\bar{\epsilon}^2}{\lambda\underline{\epsilon}} + \frac{\bar{v}^2\|\mathbf{a}_k\|_2^2\bar{\epsilon}^2}{\lambda^2\underline{\epsilon}^2}. \quad \square$$

This concludes the proof. □

Proof of Lemma A.5. Let each advertiser $i \neq k$ follow the adaptive pacing strategy. Denote $\bar{\Psi}_k(\boldsymbol{\mu}) := \Psi_k(\boldsymbol{\mu})/T$, and by $\boldsymbol{\mu}_t \in \mathbb{R}^K$ the (random) vector such that $\mu_{k,t} = \mu_k^*$ and $\mu_{i,t}$ is the multiplier of the adaptive pacing strategy of advertiser $i \neq k$ at time t . Based on the second-price allocation rule, the Lagrangian utility of advertiser k from the t^{th} auction can be written as:

$$\begin{aligned} u_{k,t} - \mu_k^* z_{k,t} + \rho_k &= \mathbf{1}\{d_{k,t} \leq b_{k,t}^\beta\} (v_{k,t} - (1 + \mu_k^*)d_{k,t}) + \rho_k \\ &\leq (v_{k,t} - (1 + \mu_k^*)d_{k,t})^+ + \rho_k, \end{aligned}$$

where the first equality follows because $z_{k,t} = \mathbf{1}\{d_{k,t} \leq b_{k,t}^\beta\} d_{k,t}$ and $u_{k,t} = \mathbf{1}\{d_{k,t} \leq b_{k,t}^\beta\} (v_{k,t} - d_{k,t})$, and the inequality because $x \leq x^+$ for all $x \in \mathbb{R}$ and dropping the indicator that advertiser k wins the auction. Taking expectations we obtain:

$$\mathbb{E}[u_{k,t} - \mu_k^* z_{k,t} + \rho_k] \stackrel{(a)}{\leq} \mathbb{E}\left[\mathbb{E}\left[(v_{k,t} - (1 + \mu_k^*)d_{k,t})^+ \mid \boldsymbol{\mu}_t\right] + \rho_k\right] \stackrel{(b)}{=} \mathbb{E}[\bar{\Psi}_k(\boldsymbol{\mu}_t)],$$

where (a) follows from conditioning on $\boldsymbol{\mu}_t$, and (b) holds since $\{v_{k,t}\}_{k=1}^K$ are independent of the multipliers $\boldsymbol{\mu}_t$. Since the dual objective is Lipschitz continuous (Lemma C.2, part (ii)) one has:

$$\bar{\Psi}_k(\boldsymbol{\mu}_t) \leq \bar{\Psi}_k(\boldsymbol{\mu}^*) + \bar{v}_k \sum_{i \neq k} a_{k,i} |\mu_{i,t} - \mu_i^*| \leq \bar{\Psi}_k(\boldsymbol{\mu}^*) + \bar{v}_k \|\mathbf{a}_k\|_2 \cdot \|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2,$$

where the last inequality follows from Cauchy-Schwarz inequality, for all $t = 1, \dots, T$. The result follows by taking expectations and using Jensen's Inequality. □

C Additional auxiliary analysis

The following result provides sufficient conditions for the dual function to be thrice differentiable and strongly convex. These conditions are required by Theorem 3.4 to show the asymptotic optimality of an adaptive pacing strategy in stationary settings. Here we assume that values and competing bids (v_k, d_k) are independently drawn from some stationary distribution. Recall that the dual function is $\Psi_k(\mu) := T(\mathbb{E}_{v_k, d_k} [(v_k - (1 + \mu)d_k)^+] + \mu\rho_k)$ and the expenditure function is $G_k(\mu) := \mathbb{E}_{v_k, d_k} [\mathbf{1}\{(1 + \mu)d_k \leq v_k\}d_k]$.

Lemma C.1. *Suppose that (a) the valuation v_k has a density satisfying $0 < \underline{f} \leq f_k(x) \leq \bar{f} < \infty$ for all $x \in [0, \bar{v}_k]$, (b) the valuation density is differentiable with bounded derivative $|f'_k(x)| \leq \bar{f}'$ for all $x \in [0, \bar{v}_k]$, (c) the competing bid d_k is independent of v_k and absolute continuous with density $h_k(x)$ satisfying $0 < \underline{h} \leq h_k(x) \leq \bar{h} < \infty$ for all $x \in [0, \bar{v}_k]$. Then:*

(i) $\Psi_k(\mu)$ is differentiable with derivative $\Psi'_k(\mu) = T(\rho_k - G_k(\mu))$.

(ii) $G_k(\mu)$ is differentiable and there exist $\underline{G}' > 0$ and $\lambda > 0$ such that the derivative is bounded by $-\underline{G}' \leq G'_k(\mu) \leq -\lambda_k < 0$ for all $\mu \in [0, \bar{\mu}_k]$. This implies that $\Psi_k(\mu)$ is strongly convex with parameter λ_k .

(iii) $G_k(\mu)$ is twice-differentiable and there exists $\bar{G}'' > 0$ such that the second derivative is upper bounded by $|G''_k(\mu)| \leq \bar{G}''$ for all $\mu \in [0, \bar{\mu}_k]$.

Proof of Lemma C.1. We drop the dependence on k to simplify the notation. Part (i) follows from Lemma C.2, part (ii). We prove the other parts of the Lemma.

(ii). From Lemma C.2, part (iii) we know that G is differentiable with derivative given in (B-2). Moreover, we have that $G'(\mu) \geq -\bar{v}^2 \bar{f}$. We can upper bound the derivative as follows:

$$G'(\mu) = - \int_0^{\bar{v}/(1+\mu)} x^2 f((1+\mu)x) dH(x) \leq -\underline{f}\underline{h} \int_0^{\bar{v}/(1+\bar{\mu})} x^2 dx = -\frac{\underline{f}\underline{h}\bar{v}^3}{3(1+\bar{\mu})^3},$$

where the first inequality follows because the densities are lower bounded, and using that $\mu \leq \bar{\mu}$ together with the fact that the integrand is positive.

(iii). Using (B-2) we obtain that the second derivative of the expenditure function is

$$G''(\mu) = \frac{\bar{v}^3}{(1+\mu)^4} f(\bar{v})h(\bar{v}/(1+\mu)) - \int_0^{\bar{v}/(1+\mu)} x^3 f'((1+\mu)x) dH(x),$$

where we used Leibniz rule because $x^2 f((1 + \mu)x)$ is differentiable almost everywhere w.r.t. μ with derivative that is bounded by $\bar{v}^3 \bar{f}'$. A similar argument as in part (ii) yields that

$$|G''(\mu)| \leq \bar{v}^3 \bar{f} \bar{h} + \frac{\bar{f}' \bar{h} \bar{v}^4}{4}. \quad \square$$

The following result obtains some key characteristics of the model primitives under the matching model described in Section 5. We denote by $a_{k,i} = \mathbb{P}\{m_{k,t} = m_{i,t}\} = \sum_{m=1}^M \alpha_{k,m} \alpha_{i,m}$ the probability that advertisers k and i compete in the same auction at a given time period. Given a vector of multipliers $\boldsymbol{\mu}$, we denote $\Psi_k(\boldsymbol{\mu}) := T(\mathbb{E}_{\mathbf{v}, \mathbf{m}}[(v_k - (1 + \mu_k)d_k)^+] + \mu_k \rho_k)$ the dual performance under $\boldsymbol{\mu}$ with $d_k = \max\left\{\max_{i: i \neq k} \{\mathbf{1}\{m_i = m_k\} v_i / (1 + \mu_i)\}, 0\right\}$ and the expectation taken w.r.t. v_i and m_i for all i . In addition, we denote by $G_k(\boldsymbol{\mu}) := \mathbb{E}_{\mathbf{v}, \mathbf{m}}[\mathbf{1}\{(1 + \mu_k)d_k \leq v_k\} d_k]$ the expected expenditure under the second-price auction allocation rule. Results for the original model hold by putting $M = 1$ and $\alpha_{k,m} = 1$.

Lemma C.2. *Suppose that for each advertiser k the valuation density satisfies $f_k(x) \leq \bar{f} < \infty$ for all $x \in [0, \bar{v}_k]$. Then:*

- (i) Fix $\boldsymbol{\mu}_{-k,t} = (\mu_{i,t})_{i \neq k} \in \mathbb{R}^{K-1}$, the competitor's multipliers at time t . The maximum competing bid $d_{k,t} = \max_{i \neq k, m_{k,t} = m_{i,t}} \{v_{i,t} / (1 + \mu_{i,t})\}$ is integrable over $[0, \bar{d}_{k,t}]$ where $\bar{d}_{k,t} = \max_{i \neq k} \{\bar{v}_i / (1 + \mu_{i,t})\}$, with cumulative distribution function

$$H_k(x; \boldsymbol{\mu}_{-k,t}) = \prod_{i \neq k} (1 - a_{k,i} \bar{F}_i((1 + \mu_{i,t})x)).$$

- (ii) $\Psi_k(\cdot)$ is differentiable and Lipschitz continuous. In particular, for any $\boldsymbol{\mu} \in \mathcal{U}$ and $\boldsymbol{\mu}' \in \mathcal{U}$, one has

$$\frac{1}{T} |\Psi_k(\boldsymbol{\mu}) - \Psi_k(\boldsymbol{\mu}')| \leq \bar{v}_k |\mu_k - \mu'_k| + \bar{v}_k \sum_{i \neq k} a_{k,i} |\mu_i - \mu'_i|.$$

- (iii) $G_k(\cdot)$ is Lipschitz continuous. In particular, for any $\boldsymbol{\mu} \in \mathcal{U}$ and $\boldsymbol{\mu}' \in \mathcal{U}$, one has

$$|G_k(\boldsymbol{\mu}) - G_k(\boldsymbol{\mu}')| \leq \bar{v}_k^2 \bar{f} |\mu_k - \mu'_k| + 2\bar{v}_k^2 \bar{f} \sum_{i \neq k} a_{k,i} |\mu_i - \mu'_i|.$$

Proof of Lemma C.2. We prove the three parts of the Lemma.

- (i). Using that the values $v_{i,t}$ are independent across advertisers and identical across time, we can write the cumulative distribution function as

$$\begin{aligned}
H_k(x; \boldsymbol{\mu}_{-k,t}) &= \mathbb{P} \left\{ \max_{i \neq k, m_{k,t} = m_{i,t}} \frac{v_{i,t}}{1 + \mu_{i,t}} \leq x \right\} \\
&\stackrel{(a)}{=} \prod_{i \neq k} \left(\mathbb{P} \{m_{k,t} \neq m_{i,t}\} + \mathbb{P} \{m_{k,t} = m_{i,t}\} \mathbb{P} \left\{ \frac{v_i}{1 + \mu_{i,t}} \leq x \right\} \right) \\
&\stackrel{(b)}{=} \prod_{i \neq k} (1 - a_{k,i} \bar{F}_i((1 + \mu_{i,t})x)) ,
\end{aligned}$$

where (a) follows by conditioning on whether advertiser $i \neq k$ participates in the same auction that advertiser k , and (b) follows from $a_{k,i} = \mathbb{P} \{m_{k,t} = m_{i,t}\}$. Because $F_i(\cdot)$ has support $[0, \bar{v}_i]$ we conclude that the support is $[0, \bar{d}_{k,t}]$ with $\bar{d}_{k,t} = \max_{i \neq k} \bar{v}_i / (1 + \mu_{i,t})$.

(ii). Denote by $\bar{\Psi}_k(\boldsymbol{\mu}) = \Psi_k(\boldsymbol{\mu})/T$. For every realized vectors $\mathbf{v} = \{v_i\}_i$ and $\mathbf{m} = \{m_i\}_i$, the function $(v_k - (1 + \mu_k)d_k)^+$ is differentiable in μ_k with derivative $-d_k \mathbf{1}\{v_k \geq (1 + \mu_k)d_k\}$, except in the set $\{(\mathbf{v}, \mathbf{m}) : v_k = (1 + \mu_k)d_k\}$ that has measure zero because values are absolutely continuous with support $[0, \bar{v}_k]$ and independent. As the derivative is bounded by d_k , which is integrable since $d_k \leq \bar{v}$ with $\bar{v} := \max_k \bar{v}_k$ from Item (i), we conclude by Leibniz's integral rule that:

$$\frac{\partial \bar{\Psi}_k(\boldsymbol{\mu})}{\partial \mu_k} = \rho_k - \mathbb{E} [d_k \mathbf{1}\{v_k \geq (1 + \mu_k)d_k\}] = \rho_k - G_k(\boldsymbol{\mu}) ,$$

which implies that $|\frac{\partial \bar{\Psi}_k(\boldsymbol{\mu})}{\partial \mu_k}| \leq \bar{v}_k$ because $\rho_k, G_k(\boldsymbol{\mu}) \in [0, \bar{v}_k]$.

Fix an advertiser $i \neq k$. Recall that the maximum competing bid faced by advertiser k is $d_k = \max_{i \neq k, m_k = m_i} \{v_i / (1 + \mu_i)\}$. Let $d_{k \setminus i} = \max_{j \neq k, i; m_k = m_j} \{v_j / (1 + \mu_j)\}$ be the maximum competing bid faced by advertiser k with advertiser i excluded. By conditioning on whether advertiser i and k compete in the same auction, we can write the random function $(v_k - (1 + \mu_k)d_k)^+$ as

$$(v_k - (1 + \mu_k)d_k)^+ = \begin{cases} \left(v_k - (1 + \mu_k)d_{k \setminus i} \vee \frac{v_i}{1 + \mu_i} \right)^+ , & \text{if } m_{k,t} = m_{i,t} , \\ \left(v_k - (1 + \mu_k)d_{k \setminus i} \right)^+ , & \text{if } m_{k,t} \neq m_{i,t} . \end{cases}$$

We obtain that the function $(v_k - (1 + \mu_k)d_k)^+$ is differentiable in μ_i , with derivative

$$v_i \frac{1 + \mu_k}{(1 + \mu_i)^2} \mathbf{1} \left\{ \frac{v_k}{1 + \mu_k} \geq \frac{v_i}{1 + \mu_i} \geq d_{k \setminus i}, m_k = m_i \right\}$$

except in the sets $\{(\mathbf{v}, \mathbf{m}) : \frac{v_k}{1 + \mu_k} = \frac{v_i}{1 + \mu_i} \geq d_{k \setminus i}, m_k = m_i\}$ and $\{(\mathbf{v}, \mathbf{m}) : \frac{v_k}{1 + \mu_k} \geq \frac{v_i}{1 + \mu_i} = d_{k \setminus i}, m_k = m_i\}$. Again, these sets have measure zero because values are absolutely continuous with support $[0, \bar{v}]$ and independent. Because the derivative is bounded by $v_k / (1 + \mu_i)$, which is integrable since $v_k \leq \bar{v}$, we conclude by Leibniz's integral rule that:

$$\begin{aligned}\frac{\partial \bar{\Psi}_k(\boldsymbol{\mu})}{\partial \mu_i} &= \mathbb{E} \left[\frac{v_i}{1 + \mu_i} \frac{1 + \mu_k}{(1 + \mu_i)^2} \mathbf{1} \left\{ \frac{v_k}{1 + \mu_k} \geq \frac{v_i}{1 + \mu_i} \geq d_{k \setminus i}, m_k = m_i \right\} \right] \\ &\leq \mathbb{E} \left[\frac{v_k}{1 + \mu_i} \mathbf{1} \{m_k = m_i\} \right] \leq \bar{v}_k a_{k,i},\end{aligned}$$

where the first inequality follows because $\frac{v_i}{1 + \mu_i} \leq \frac{v_k}{1 + \mu_k}$ and dropping part of the indicator, and the last inequality follows because $v_k \in [0, \bar{v}_k]$ and $\mu_i \geq 0$. This concludes the proof.

(iii). We show that $G_k(\boldsymbol{\mu})$ is Lipschitz continuous by bounding its derivatives. Since values are independent across advertisers we can write the expected expenditure as

$$\begin{aligned}G_k(\boldsymbol{\mu}) &= \int_0^{\bar{v}_k} x \bar{F}_k((1 + \mu_k)x) dH_k(x; \boldsymbol{\mu}_{-k}) \\ &= \int_0^{\bar{v}_k} ((1 + \mu_k)x f_k((1 + \mu_k)x) - \bar{F}_k((1 + \mu_k)x)) H_k(x; \boldsymbol{\mu}_{-k}) dx,\end{aligned}$$

where the second equation follows from integration by parts. Using the first expression for the expected expenditure we obtain that

$$\frac{\partial G_k(\boldsymbol{\mu})}{\partial \mu_k} = - \int_0^{\bar{v}_k} x^2 f_k((1 + \mu_k)x) \mathbf{1} \{(1 + \mu_k)x \leq \bar{v}_k\} dH_k(x; \boldsymbol{\mu}_{-k}), \quad (\text{B-2})$$

where we used Leibniz rule because $x \bar{F}_k((1 + \mu_k)x)$ is differentiable w.r.t. μ_k almost everywhere with derivative that is bounded by $\bar{v}_k^2 \bar{f}$. Therefore, one obtains

$$\left| \frac{\partial G_k(\boldsymbol{\mu})}{\partial \mu_k} \right| \leq \bar{v}_k^2 \bar{f}.$$

Using the second expression for the expected expenditure we obtain for $i \neq k$ that

$$\frac{\partial G_k(\boldsymbol{\mu})}{\partial \mu_i} = \int_0^{\bar{v}_k} ((1 + \mu_k)x f_k((1 + \mu_k)x) - \bar{F}_k((1 + \mu_k)x)) \frac{\partial H_k}{\partial \mu_i}(x; \boldsymbol{\mu}_{-k}) dx, \quad (\text{B-3})$$

where we used Leibniz rule as $H_k(x; \boldsymbol{\mu}_{-k})$ is differentiable w.r.t. μ_i almost everywhere with derivative

$$\frac{\partial H_k}{\partial \mu_i}(x; \boldsymbol{\mu}_{-k}) = a_{k,i} x f_i((1 + \mu_i)x) \mathbf{1} \{(1 + \mu_i)x \leq \bar{v}_i\} \prod_{j \neq k, i} (1 - a_{k,j} \bar{F}_j((1 + \mu_j)x)),$$

which is bounded from above by $a_{k,i} \bar{v}_k \bar{f} \mathbf{1} \{(1 + \mu_i)x \leq \bar{v}_i\}$. Therefore, one obtains

$$\left| \frac{\partial G_k(\boldsymbol{\mu})}{\partial \mu_i} \right| \leq 2a_{k,i} \bar{v}_k^2 \bar{f},$$

and the result follows. \square

C.1 Auxiliary results

Proposition C.3. (Uniqueness of steady state) *Suppose Assumption 4.1 holds. Then there exists a unique vector of multipliers $\boldsymbol{\mu}^* \in \mathcal{U}$ defined by (5).*

Proof. We first establish that selecting a multiplier outside of $[0, \bar{\mu}_k]$ is a dominated strategy for advertiser k . Notice that for every $\boldsymbol{\mu}_{-k}$ and $x > \bar{\mu}_k$ we have that

$$\Psi_k(x, \boldsymbol{\mu}_{-k}) = \mathbb{E}_{\mathbf{v}} \left[(v_k - (1+x)d_k)^+ \right] + x\rho_k \stackrel{(a)}{\geq} x\rho_k \stackrel{(b)}{>} \bar{v} \stackrel{(c)}{\geq} \Psi_k(0, \boldsymbol{\mu}_{-k}),$$

where (a) follows from dropping the first term, (b) holds since by Assumption 4.1 one has that $\rho_k \geq \bar{v}_k/\bar{\mu}_k$ from and $x > \bar{\mu}_k$, and (c) follows from $0 \leq v_{k,t} \leq \bar{v}_k$. Thus every $x > \bar{\mu}_k$ in the dual problem is dominated by $x = 0$, and the equilibrium multipliers lie in the set \mathcal{U} . Define $G_k(\boldsymbol{\mu}) := \mathbb{E}_{\mathbf{v}} [\mathbf{1}\{(1 + \mu_k)d_k \leq v_k\}d_k]$ to be the expected expenditure under the second-price auction allocation rule. Assumption 4.1 implies that:

$$(\boldsymbol{\mu} - \boldsymbol{\mu}^*)^\top (\mathbf{G}(\boldsymbol{\mu}^*) - \mathbf{G}(\boldsymbol{\mu})) > 0, \tag{B-4}$$

for all $\boldsymbol{\mu} \in \mathcal{U}$ such that $\boldsymbol{\mu} \neq \boldsymbol{\mu}^*$. To prove uniqueness, suppose that there exists another equilibrium multiplier $\boldsymbol{\mu} \in \mathcal{U}$ such that $\boldsymbol{\mu} \neq \boldsymbol{\mu}^*$. From (B-4) one has:

$$\begin{aligned} 0 &< \sum_{k=1}^K (\mu_k - \mu_k^*) (G_k(\boldsymbol{\mu}^*) - \rho_k + \rho_k - G_k(\boldsymbol{\mu})) \\ &\stackrel{(a)}{=} \sum_{k=1}^K \mu_k (G_k(\boldsymbol{\mu}^*) - \rho_k) - \mu_k^* (\rho_k - G_k(\boldsymbol{\mu})), \end{aligned}$$

where (a) follows from $\mu_k(\rho_k - G_k(\boldsymbol{\mu})) = 0$ and $\mu_k^*(G_k(\boldsymbol{\mu}^*) - \rho_k) = 0$ by (5). As $\mu_k, \mu_k^* \geq 0$ and $G_k(\boldsymbol{\mu}^*), G_k(\boldsymbol{\mu}) \leq \rho_k$, we obtain that the right hand-side is non-positive, contradicting (B-4). \square

Lemma C.4. *Let $\{\delta_t\}_{t \geq 1}$ be a sequence of numbers such that $\delta_t \geq 0$ and $\delta_{t+1} \leq (1-a)\delta_t + b_t$ with $b_t \geq 0$ and $0 \leq a \leq 1$. Then,*

$$\delta_t \leq (1-a)^{t-1} \delta_1 + \sum_{s=1}^{t-1} (1-a)^{t-1-s} b_s.$$

When $b_t = b$ for all $t \geq 1$ and $a > 0$, we obtain

$$\delta_t \leq (1-a)^{t-1} \delta_1 + \frac{b}{a}.$$

Proof. We prove the result by induction. The result trivially holds for $t = 1$ because $a, b_1 \geq 0$. For $t > 1$, the recursion gives

$$\delta_{t+1} \leq (1-a)\delta_t + b_t \leq (1-a)^t \delta_1 + \sum_{s=1}^t (1-a)^{t-s} b_s,$$

where the second inequality follows from the induction hypothesis and the fact that $1-a \geq 0$. The last inequality follows because $\sum_{s=1}^{t-1} (1-a)^{t-1-s} \leq \sum_{s=0}^{\infty} (1-a)^s = 1/a$ when $a \in (0, 1]$. \square

Lemma C.5. *Let $\{\delta_t\}_{t \geq 1}$ be a sequence of numbers such that $\delta_t \geq 0$ and $\delta_{t+1} \leq (1-a)\delta_t + b + c\delta_t^{1/2}$ with $c \geq 0$, $b \geq 0$ and $0 \leq a \leq 1$. Then,*

$$\delta_t \leq (1-a/2)^{t-1} \delta_1 + \frac{2b}{a} + \frac{c^2}{a^2}.$$

Proof. The square root term can be bounded as follows

$$c\delta_t^{1/2} = \frac{c}{a^{1/2}} a^{1/2} \delta_t^{1/2} \leq \frac{c^2}{2a} + \frac{a\delta_t}{2},$$

because $xy \leq (x^2 + y^2)/2$ for $x, y \in \mathbb{R}$ by the AM-GM inequality. Using this bound, we can rewrite the inequality in the statement as

$$\delta_{t+1} \leq (1-a)\delta_t + b + c\delta_t^{1/2} \leq (1-a/2)\delta_t + b + c^2/2a.$$

The result then follows from Lemma C.4. \square

C.2 Stability analysis

We first show that the first part of Assumption 4.1 can be implied by the diagonal strict concavity condition defined in Rosen (1965). Indeed, since the set \mathcal{U} is compact and since the vector function $\mathbf{G}(\cdot)$ is bounded in \mathcal{U} , to verify that the first part of Assumption 4.1 holds, it suffices to show that $(\boldsymbol{\mu} - \boldsymbol{\mu}')^\top (\mathbf{G}(\boldsymbol{\mu}') - \mathbf{G}(\boldsymbol{\mu})) > 0$ for all $\boldsymbol{\mu}, \boldsymbol{\mu}' \in \mathcal{U}$. The latter is equivalent to the diagonal strict concavity assumption of Rosen (1965). Furthermore, denote by $J_{\mathbf{G}} : \mathbb{R}_+^K \rightarrow \mathbb{R}^{K \times K}$ the Jacobian matrix of the vector function \mathbf{G} , that is, $J_{\mathbf{G}}(\boldsymbol{\mu}) = \left(\frac{\partial G_k}{\partial \mu_i}(\boldsymbol{\mu}) \right)_{k,i}$. Then, by Theorem 6 of Rosen (1965), it is sufficient to show that the symmetric matrix $J_{\mathbf{G}}(\boldsymbol{\mu}) + J_{\mathbf{G}}^\top(\boldsymbol{\mu})$ is negative definite.

We next provide analytical expressions for $\mathbf{G}(\boldsymbol{\mu})$ for two advertisers with valuations that are independently uniformly distributed and exponentially distributed to demonstrate numerically that the latter condition holds in these cases.

Example C.6. (Two bidders with uniform valuations) Assume $K = 2$, $\mathcal{U} = [0, 2]^2$, and $v_{k,t} \sim U[0, 1]$, i.i.d. for all $k \in \{1, 2\}$ and $t \in \{1, \dots, T\}$. One obtains:

$$\begin{aligned} \mathbf{G}_1(\boldsymbol{\mu}) &= \int_0^1 \int_0^{\min\{\frac{1+\mu_2}{1+\mu_1}x_1, 1\}} \frac{x_2}{1+\mu_2} dx_2 dx_1 \\ &= \mathbf{1}\{\mu_2 \leq \mu_1\} \frac{1+\mu_2}{6(1+\mu_1)^2} + \mathbf{1}\{\mu_1 < \mu_2\} \left(\frac{1}{2(1+\mu_2)} - \frac{1+\mu_1}{3(1+\mu_2)^2} \right). \end{aligned}$$

Therefore:

$$\mathbf{G}(\boldsymbol{\mu}) = \mathbf{1}\{\mu_2 \leq \mu_1\} \begin{bmatrix} \frac{1+\mu_2}{6(1+\mu_1)^2} \\ \frac{1}{2(1+\mu_1)} - \frac{1+\mu_2}{3(1+\mu_1)^2} \end{bmatrix} + \mathbf{1}\{\mu_1 < \mu_2\} \begin{bmatrix} \frac{1}{2(1+\mu_2)} - \frac{1+\mu_1}{3(1+\mu_2)^2} \\ \frac{1+\mu_1}{6(1+\mu_2)^2} \end{bmatrix}$$

Following this expression, one may validate the first part of Assumption 4.1 by creating a grid of $\boldsymbol{\mu}_{i,j} \in \mathcal{U}$, and for a given grid calculate the maximal monotonicity constant λ for which the condition holds. For example, for a 20×20 grid (with $\|\boldsymbol{\mu}_{i,j} - \boldsymbol{\mu}_{i,j+1}\| = \|\boldsymbol{\mu}_{i,j} - \boldsymbol{\mu}_{i+1,j}\| = 0.1$ for all $i = 0, 1, \dots, 19$ and $j = 0, 1, \dots, 19$) the latter condition holds with $\lambda = 0.013$. Similarly, for a 40×40 grid (with $\|\boldsymbol{\mu}_{i,j} - \boldsymbol{\mu}_{i,j+1}\| = \|\boldsymbol{\mu}_{i,j} - \boldsymbol{\mu}_{i+1,j}\| = 0.05$ for all $i = 0, 1, \dots, 39$ and $j = 0, 1, \dots, 39$) the latter condition holds with $\lambda = 0.0127$. \square

Example C.7. (Two bidders with exponential valuations) Assume $K = 2$, $\mathcal{U} = [0, 2]^2$, and $v_{k,t} \sim \exp(1)$, i.i.d. for all $k \in \{1, 2\}$ and $t \in \{1, \dots, T\}$. One obtains:

$$\mathbf{G}_1(\boldsymbol{\mu}) = \int_0^\infty \int_0^{\frac{1+\mu_2}{1+\mu_1}x_1} \frac{x_2 e^{-x_2} e^{-x_1}}{1+\mu_2} dx_2 dx_1 = \frac{1+\mu_2}{(2+\mu_2+\mu_1)^2}.$$

Following same lines as in Example C.7, one may validate the first part of Assumption 4.1. For example, for a 20×20 grid (with $\|\boldsymbol{\mu}_{i,j} - \boldsymbol{\mu}_{i,j+1}\| = \|\boldsymbol{\mu}_{i,j} - \boldsymbol{\mu}_{i+1,j}\| = 0.1$ for all $i = 0, 1, \dots, 19$ and $j = 0, 1, \dots, 19$) the latter condition holds with $\lambda = 0.0295$. Similarly, for a 40×40 grid (with $\|\boldsymbol{\mu}_{i,j} - \boldsymbol{\mu}_{i,j+1}\| = \|\boldsymbol{\mu}_{i,j} - \boldsymbol{\mu}_{i+1,j}\| = 0.05$ for all $i = 0, 1, \dots, 39$ and $j = 0, 1, \dots, 39$) the latter condition holds with $\lambda = 0.0286$. \square

The following result expands and complements the above examples by showing that when the number of players is large, the stability assumption holds in symmetric settings in which every advertiser participates in each auction with the same probability and all advertisers have the same distribution of values. In particular, the monotonicity constant λ of the expenditure function \mathbf{G} is shown to be independent of the number of players.

Proposition C.8. (Stability in symmetric settings) *Consider a symmetric setting in which advertisers participate in each auction with probability $\alpha_{i,m} = 1/M$, advertisers values are drawn from a continuous density $f(\cdot)$, and the ratio of number of auctions to number of players $\kappa := K/M$ is fixed. Then, there exist $\underline{K} \in \mathbb{N}$ and $\lambda > 0$ such that for all $K \geq \underline{K}$ there exists a set $\mathcal{U} \subset \mathbb{R}_+^K$ with $\mathbf{0} \in \mathcal{U}$ such that \mathbf{G} is λ -strongly monotone over \mathcal{U} .*

Proof. We prove the result in three steps. First, we argue that a sufficient condition for λ -strong monotonicity of \mathbf{G} over \mathcal{U} is that λ is a lower bound on the minimum eigenvalue of the symmetric part of the Jacobian of $-\mathbf{G}$ over \mathcal{U} . Second, we characterize the Jacobian of the vector function \mathbf{G} in the general case. Third, we show that in the symmetric case \mathbf{G} is strongly monotone around $\boldsymbol{\mu} = \mathbf{0}$ by bounding the minimum eigenvalue of the symmetric part of the Jacobian of $-\mathbf{G}$ at zero.

Step 1. We denote by $J_{\mathbf{G}} : \mathbb{R}_+^K \rightarrow \mathbb{R}^{K \times K}$ the Jacobian matrix of the vector function \mathbf{G} , that is, $J_{\mathbf{G}}(\boldsymbol{\mu}) = \left(\frac{\partial G_k}{\partial \mu_i}(\boldsymbol{\mu}) \right)_{k,i}$. Let λ be a lower bound on the minimum eigenvalue of the symmetric part of $-J_{\mathbf{G}}(\boldsymbol{\mu})$ over \mathcal{U} . That is, λ satisfies

$$\lambda \leq \min_{\|\mathbf{x}\|_2=1} -\frac{1}{2} \mathbf{x}^\top (J_{\mathbf{G}}(\boldsymbol{\mu}) + J_{\mathbf{G}}(\boldsymbol{\mu})^\top) \mathbf{x}$$

for all $\boldsymbol{\mu} \in \mathcal{U}$. Lemma C.2, part (iii) shows that $\mathbf{G}(\boldsymbol{\mu})$ is differentiable in $\boldsymbol{\mu}$. Thus, by the mean value theorem there exists some $\boldsymbol{\xi} \in \mathbb{R}^K$ in the segment between $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$ such that for all $\boldsymbol{\mu}, \boldsymbol{\mu}' \in \mathcal{U}$ one has $\mathbf{G}(\boldsymbol{\mu}') = \mathbf{G}(\boldsymbol{\mu}) + J_{\mathbf{G}}(\boldsymbol{\xi})(\boldsymbol{\mu}' - \boldsymbol{\mu})$. Therefore,

$$(\boldsymbol{\mu} - \boldsymbol{\mu}')^\top (\mathbf{G}(\boldsymbol{\mu}') - \mathbf{G}(\boldsymbol{\mu})) = -\frac{1}{2} (\boldsymbol{\mu} - \boldsymbol{\mu}')^\top (J_{\mathbf{G}}(\boldsymbol{\xi}) + J_{\mathbf{G}}(\boldsymbol{\xi})^\top) (\boldsymbol{\mu} - \boldsymbol{\mu}') \geq \lambda \|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_2^2,$$

since $\boldsymbol{\xi} \in \mathcal{U}$. Hence, a sufficient condition for λ -strong monotonicity of \mathbf{G} over \mathcal{U} is that λ is a lower bound on the minimum eigenvalue of the symmetric part of the Jacobian of $-\mathbf{G}$ over \mathcal{U} .

Step 2. Some definitions are in order. Let $H_{k,i}(x; \boldsymbol{\mu}_{-k,i}) = \prod_{j \neq i,k} (1 - a_{k,j} \bar{F}_j((1 + \mu_j)x))$ and $\ell_k(x) = x f_k(x) \mathbf{1}\{x \leq \bar{v}_k\}$. Additionally, we denote

$$\begin{aligned} \gamma_{k,i} &= \int_0^{\bar{v}_k} \ell_k((1 + \mu_k)x) \ell_i((1 + \mu_i)x) H_{k,i}(x; \boldsymbol{\mu}_{-k,i}) dx, \\ \omega_{k,i} &= \int_0^{\bar{v}_k} \bar{F}_k((1 + \mu_k)x) \ell_i((1 + \mu_i)x) H_{k,i}(x; \boldsymbol{\mu}_{-k,i}) dx. \end{aligned}$$

We first determine the partial derivatives of the cumulative distribution function H_k . One has

$$\begin{aligned}\frac{\partial H_k}{\partial x}(x; \boldsymbol{\mu}_{-k}) &= \sum_{i \neq k} a_{k,i} (1 + \mu_i) f_i((1 + \mu_i)x) \mathbf{1}\{(1 + \mu_i)x \leq \bar{v}\} H_{k,i}(x; \boldsymbol{\mu}_{-k,i}), \\ \frac{\partial H_k}{\partial \mu_i}(x; \boldsymbol{\mu}_{-k}) &= a_{k,i} x f_i((1 + \mu_i)x) \mathbf{1}\{(1 + \mu_i)x \leq \bar{v}\} H_{k,i}(x; \boldsymbol{\mu}_{-k,i}).\end{aligned}$$

Using equation (B-2) one obtains that

$$\begin{aligned}\frac{\partial G_k}{\partial \mu_k}(\boldsymbol{\mu}) &= - \int_0^{\bar{v}_k} x^2 f_k((1 + \mu_k)x) \mathbf{1}\{(1 + \mu_k)x \leq \bar{v}\} dH_k(x; \boldsymbol{\mu}_{-k}) \\ &= - \sum_{i \neq k} \frac{a_{k,i}}{1 + \mu_k} \int_0^{\bar{v}_k} \ell_k((1 + \mu_k)x) \ell_i((1 + \mu_i)x) H_{k,i}(x; \boldsymbol{\mu}_{-k,i}) dx \\ &= - \sum_{i \neq k} \frac{a_{k,i}}{1 + \mu_k} \gamma_{k,i}.\end{aligned}$$

Using equation (B-3) we have

$$\begin{aligned}\frac{\partial G_k}{\partial \mu_i}(\boldsymbol{\mu}) &= \int_0^{\bar{v}_k} ((1 + \mu_k)x f_k((1 + \mu_k)x) - \bar{F}_k((1 + \mu_k)x)) \frac{\partial H_k}{\partial \mu_i}(x; \boldsymbol{\mu}_{-k}) dx \\ &= \frac{a_{k,i}}{1 + \mu_i} \int_0^{\bar{v}_k} \ell_k((1 + \mu_k)x) \ell_i((1 + \mu_i)x) H_{k,i}(x; \boldsymbol{\mu}_{-k,i}) dx \\ &\quad - \frac{a_{k,i}}{1 + \mu_i} \int_0^{\bar{v}_k} \bar{F}_k((1 + \mu_k)x) \ell_i((1 + \mu_i)x) H_{k,i}(x; \boldsymbol{\mu}_{-k,i}) dx \\ &= \frac{a_{k,i}}{1 + \mu_i} (\gamma_{k,i} - \omega_{k,i}).\end{aligned}$$

Step 3. Consider a symmetric setting in which each advertiser participates in each auction with the same probability and all advertisers have the same distribution of values. By symmetry we obtain that $a_{k,i} = 1/M$, because $\alpha_{i,m} = 1/M$ for all advertiser i and auction m . Evaluating at $\boldsymbol{\mu} = \mathbf{0}$ we obtain that $\gamma_{k,i} = \gamma := \int_0^{\bar{v}} \ell(x)^2 (1 - \bar{F}(x)/M)^{K-2} dx$ for all $k \neq i$ and $\omega_{k,i} = \omega := \int_0^{\bar{v}} \bar{F}(x) \ell(x) (1 - \bar{F}(x)/M)^{K-2} dx$ for all $k \neq i$. Therefore, $\frac{\partial G_k}{\partial \mu_k}(\mathbf{0}) = -\frac{K-1}{M} \gamma$ and $\frac{\partial G_k}{\partial \mu_i}(\mathbf{0}) = \frac{1}{M} (\gamma - \omega)$. The eigenvalues of $-(J_G(\mathbf{0}) + J_G(\mathbf{0})^\top)/2$ are $\nu_1 = \frac{K-1}{M} \omega$ with multiplicity 1 and $\nu_2 = \frac{K\gamma - \omega}{M}$ with multiplicity $K - 1$. Assume further that the expected number of players per auction $\kappa := K/M$ is fixed, which implies that the number of auctions is proportional to the number of players. Because $(1 - \bar{F}(x)/M)^{K-2}$ converges to $e^{-\kappa \bar{F}(x)}$ as $K \rightarrow \infty$ and the integrands are bounded, Dominated Convergence Theorem implies that

$$\lim_{K \rightarrow \infty} \nu_1 = \kappa \int_0^{\bar{v}} \bar{F}(x) \ell(x) e^{-\kappa \bar{F}(x)} dx > 0,$$

and

$$\lim_{K \rightarrow \infty} \nu_2 = \kappa \int_0^{\bar{v}} \ell(x)^2 e^{-\kappa \bar{F}(x)} dx > 0.$$

Hence, there exist $\underline{K} \in \mathbb{N}$ and $\lambda' > 0$ such that for all $K \geq \underline{K}$ the minimum eigenvalue value of $-(J_{\mathbf{G}}(\mathbf{0}) + J_{\mathbf{G}}(\mathbf{0})^\top)/2$ is at least $\lambda' > 0$. Because densities are continuous, one obtains that $J_{\mathbf{G}}(\boldsymbol{\mu})$ is continuous in $\boldsymbol{\mu}$. Since the eigenvalues of a matrix are continuous functions of its entries, we conclude that there exists $\lambda \in (0, \lambda']$ such that for each $K \geq \underline{K}$ there exists a set $\mathcal{U} \subset \mathbb{R}^K$ with $\mathbf{0} \in \mathcal{U}$ such that \mathbf{G} is λ -strongly monotone over \mathcal{U} . \square

D Numerical analysis of convergence under simultaneous learning

We next describe the setup and results of numerical experiments we conducted to demonstrate the convergence of dual multipliers established in Theorem 4.3 and the convergence in performance established in Theorem 4.4.

Setup and methodology. We simulate the sample path and payoff achieved when K bidders with symmetric target expenditure rates ρ follow adaptive pacing strategies throughout synchronous campaigns of T periods. (The parametric values we tested are provided below). In each period, each advertiser observed a valuation that was drawn independently across advertisers and time periods from a uniform distribution over $[0, 1]$. All the advertisers followed an adaptive pacing strategy tuned by a step size ε , and an initial dual multiplier μ_1 . The upper bound on the dual multipliers was set to $\bar{\mu} = 2/\rho$, guaranteeing $\bar{\mu} > \bar{v}/\rho$ so the second part of Assumption 4.1 holds.

Recall that $\boldsymbol{\mu}^*$ is the profile of dual multipliers that solves the complementarity conditions given in (5). We note that under uniform valuations one may obtain analytical solution of the form

$$\mu_k^* = \left[\frac{K-1}{K(K+1)\rho} - 1 \right]^+,$$

for all $k \in \{1, \dots, K\}$. Given a vector of multipliers $\boldsymbol{\mu}$, recall that one denotes

$$\Psi_k(\boldsymbol{\mu}) = T \left(\mathbb{E}_{\mathbf{v}} \left[(v_k - (1 + \mu_k)d_k)^+ \right] + \mu_k \rho_k \right),$$

with d_k as defined in §4.1. Recalling Equation (4), $\Psi_k(\boldsymbol{\mu}^*)$ is the dual performance of advertiser k when in each period t each advertiser i bids $b_{i,t} = v_{i,t}/(1 + \mu_i^*)$. Similarly to the above, we note

that under the uniform valuations one may obtain an analytical solutions of the form

$$\Psi_k(\boldsymbol{\mu}^*) = T \left(\frac{1}{K(K+1)} + \left[\frac{K-1}{K(K+1)} - \rho \right]^+ \right),$$

for all $k \in \{1, \dots, K\}$. Following a profile of adaptive pacing strategies, we generated a sample path of multipliers' profiles $\boldsymbol{\mu}_t$ and the corresponding sequences of payoffs Π_k^A , for all $k = 1, \dots, K$. Then, the time-average mean squared error of the sample path of dual multipliers is given by

$$MSE = \frac{1}{T} \sum_{t=1}^T \mathbb{E} [\|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_2^2],$$

and the average loss relative to the profile $\boldsymbol{\mu}^*$ is given by

$$L = \frac{1}{KT} \sum_{k=1}^K (\Psi_k(\boldsymbol{\mu}^*) - \Pi_k^A).$$

We explored all combinations of horizon lengths $T \in \{100, 500, 1000, 5000, 10000, 50000\}$, number of bidders $K \in \{2, 5, 10\}$, and target expenditure rates $\rho \in \{0.1, 0.3, 0.5, 0.7\}$, as well as the adaptive pacing strategy step size $\varepsilon \in \{T^{-1}, T^{-2/3}, T^{-1/2}\}$ and initial selection of dual multipliers $\mu_1 \in \{0.001, 0.01, 0.1, 0.3, 0.5, 0.7\}$. In addition to using the *original* update rule of the adaptive pacing strategy, that is, $\mu_{k,t+1} = P_{[0, \bar{\mu}_k]}(\mu_{k,t} - \epsilon_{k,t}(\rho - z_{k,t}))$, we also experimented with the following *adjusted* update rule:

$$\mu_{k,t+1} = P_{[0, \bar{\mu}_k]} \left(\mu_{k,t} - \epsilon_{k,t} \left(\rho - \frac{1}{t} \sum_{s=1}^t z_{k,s} \right) \right). \quad (\text{B-5})$$

This update rule is an example of a rule that depends on the entire history. At each time period $t + 1$, the direction of the gradient step at that time depends on the difference between the target expenditure rate ρ and the average expenditure rate up to time t . Overall, we explored 2,592 combinations of problem parameters T , K , and ρ , as well as the strategy's parameters ε and μ_1 . Each instance was replicated 100 times, leading to low mean standard errors.

Results. The results of our numerical analysis support the theoretical convergence established in Theorems 4.3 and 4.4. The key findings are highlighted through representative examples in Figure 2. The upper-left and upper-right parts of Figure 2 demonstrate the asymptotic convergence of dual multipliers and performance, respectively, for two bidders with target expenditure rate $\rho = 0.1$ and initial multiplier $\mu_1 = 0.001$. This selection of initial dual multipliers illustrates a case where bidders begin by bidding truthfully, and then learn the extent to which they need to shade bids

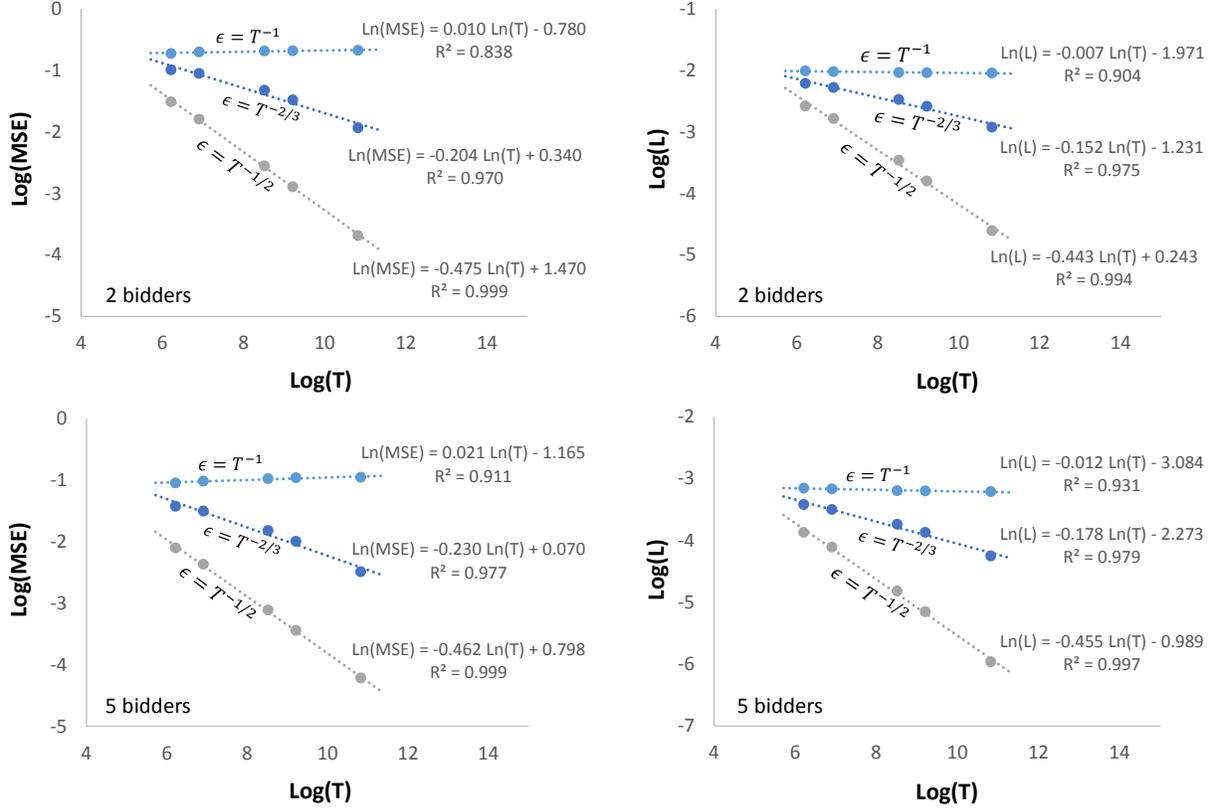


Figure 2: Convergence under simultaneous adoption of adaptive pacing strategies. (*Upper left*) The linear relations in the log-log plots illustrate the asymptotic convergence rates of dual multipliers for two symmetric bidders with target expenditure rate $\rho = 0.1$ and initial multiplier $\mu_1 = 0.001$, under different selections of step sizes. (*Upper right*) Convergence in performance for two bidders. (*Lower left*) Convergence in dual multipliers for five bidders. (*Lower right*) Convergence in performance for five bidders.

throughout the campaign (in this case $\mu_k^* = 0.66$ for $k \in \{1, 2\}$). The linear relations in the log-log plots illustrate the convergence rates under different selections of step sizes. Notably, when the step size is $\epsilon = T^{-1}$, the third part of Assumption 4.2 does not hold, and indeed the profile of dual multipliers and the average payoff do not converge to μ^* and $\Psi_k(\mu^*)$, respectively. On the other hand, under step size selections of $\epsilon = T^{-2/3}$ and $\epsilon = T^{-1/2}$, Assumption 4.2 holds. (In Example C.7, Appendix C, we demonstrate a numerical estimation of the monotonicity parameter for two bidders with valuations that are drawn from a uniform distribution over $[0, 1]$, and establish that in such a case $\lambda \sim 0.013$.) Indeed, with both of these step sizes the dual multipliers and the average payoff converge to μ^* and $\Psi_k(\mu^*)$, respectively. A step size selection of $\epsilon = T^{-1/2}$ led to superior convergence rates of $T^{-1/2}$, for both multipliers and performance. This supports the multiplier convergence rate we established theoretically when $\epsilon = T^{-1/2}$, as well as our conjecture on the performance convergence rate with the same step size.

The lower parts of Figure 2 provide similar results for the case of 5 bidders (in such a case $\mu_k^* = 0.33$ for all $k \in \{1, \dots, 5\}$). Results are consistent across different horizon lengths T , number of bidders K , target expenditure rates ρ , and initial multiplier μ_1 . Notably, when the initial multiplier selection μ_1 is higher than μ_k^* , all the bidders shade bids aggressively at early stages. This leads to lower initial payments and, in turn, overall performance that is better than the limit performance $\Psi_k(\mu^*)$. The adjusted step size in (B-5) achieved performance that is very similar to the one of the original update rule of the adaptive pacing strategy. The complete set of results for all the numerical experiments is with the authors and is available upon request.

References

- Araman, V. F. and Caldentey, R. (2011), Revenue management with incomplete demand information, in ‘Encyclopedia of Operations Research’, Wiley.
- Ashlagi, I., Monderer, D. and Tennenholtz., M. (2012), ‘Robust learning equilibrium’, *arXiv preprint arXiv:1206.6826*.
- Badanidiyuru, A., Kleinberg, R. and Slivkins, A. (2013), Bandits with knapsacks, in ‘Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science’, FOCS ’13, IEEE Computer Society, Washington, DC, USA, pp. 207–216.
- Badanidiyuru, A., Langford, J. and Slivkins, A. (2014), Resourceful contextual bandits, in M. F. Balcan, V. Feldman and C. Szepesvri, eds, ‘Proceedings of The 27th Conference on Learning Theory’, Vol. 35 of *Proceedings of Machine Learning Research*, PMLR, Barcelona, Spain, pp. 1109–1134.
- Balseiro, S., Kim, A., Mahdian, M. and Mirrokni, V. (2017), Budget management strategies in repeated auctions, in ‘22nd International Conference on World Wide Web, WWW ’17, Perth, Australia, 3-7 April, 2017’, p. To appear.
- Balseiro, S. R., Besbes, O. and Weintraub, G. Y. (2015), ‘Repeated auctions with budgets in ad exchanges: Approximations and design’, *Management Science* **61**(4), 864–884.
- Balseiro, S. R., Feldman, J., Mirrokni, V. and Muthukrishnan, S. (2014), ‘Yield optimization of display advertising with ad exchange’, *Management Science* **60**(12), 2886–2907.
- Beck, A. and Teboulle, M. (2003), ‘Mirror descent and nonlinear projected subgradient methods for convex optimization’, *Operations Research Letters* **31**(3), 167–175.
- Benveniste, A., Priouret, P. and Metivier, M. (1990), *Adaptive algorithms and stochastic approximations*, Springer-Verlag, New York.
- Besbes, O., Gur, Y. and Zeevi, A. (2015), ‘Non-stationary stochastic optimization’, *Operations research* **63**(5), 1227–1244.
- Borgs, C., Chayes, J., Immorlica, N., Jain, K., Etesami, O. and Mahdian, M. (2007), Dynamics of bid optimization in online advertisement auctions, in ‘Proceedings of the 16th International Conference on World Wide Web’, WWW ’07, ACM, New York, NY, USA, pp. 531–540.
- Borodin, A. and El-Yaniv, R. (1998), *Online Computation and Competitive Analysis*, Cambridge University Press, New York, NY, USA.

- Brafman, R. I. and Tennenholtz, M. (2004), ‘Efficient learning equilibrium’, *Artificial Intelligence* **159**, 27–47.
- Celis, L. E., Lewis, G., Mobius, M. and Nazerzadeh, H. (2014), ‘Buy-it-now or take-a-chance: Price discrimination through randomized auctions’, *Management Science* **60**(12), 2927–2948.
- Cesa-Bianchi, N. and Lugosi, G. (2006), *Prediction, Learning, and Games*, Cambridge University Press, Cambridge, UK.
- Charles, D. X., Chakrabarty, D., Chickering, M., Devanur, N. R. and Wang, L. (2013), Budget smoothing for internet ad auctions: a game theoretic approach, *in* ‘ACM Conference on Electronic Commerce, EC ’13, Philadelphia, PA, USA, June 16-20, 2013’, pp. 163–180.
- Conitzer, V., Kroer, C., Sodomka, E. and Stier-Moses, N. E. (2017), ‘Multiplicative pacing equilibria in auction markets’, *Working Paper* .
- Duchi, J., Hazan, E. and Singer, Y. (2011), ‘Adaptive subgradient methods for online learning and stochastic optimization’, *Journal of Machine Learning Research* **12**(Jul), 2121–2159.
- eMarketer (2016a), ‘More than two-thirds of us digital display ad spending is programmatic’, <https://www.emarketer.com/Article/More-Than-Two-Thirds-of-US-Digital-Display-Ad-Spending-Programmatic/1013789>. Accessed: 2017-07-10.
- eMarketer (2016b), ‘Us digital ad spending to surpass tv this year’, <https://www.emarketer.com/Article/US-Digital-Ad-Spending-Surpass-TV-this-Year/1014469>. Accessed: 2017-07-10.
- Fudenberg, D. and Levine, D. (1998), *The Theory of Learning in Games*, MIT Press, Cambridge MA.
- Han, Z., Zheng, R. and Poor, H. V. (2011), ‘Repeated auctions with bayesian nonparametric learning for spectrum access in cognitive radio networks’, *IEEE Transactions on Wireless Communications* **10**(3), 890–900.
- Hannan, J. (1957), *Approximation to bayes risk in repeated plays, Contributions to the Theory of Games, Volume 3*, Princeton University Press, Cambridge, UK.
- Hart, S. and Mas-Colell, A. (2000), ‘A simple adaptive procedure leading to correlated equilibria’, *Econometrica* **68**(5), 1127–1150.
- Hart, S. and Mas-Colell, A. (2013), *Simple Adaptive Strategies*, Vol. 4. World Scientific Publishing Company.
- Hazan, E., Agarwal, A. and Kale, S. (2007), ‘Logarithmic regret algorithms for online convex optimization’, *Mach. Learn.* **69**(2-3), 169–192.
- Hon-Snir, S., Monderer, D. and Sela, A. (1998), ‘A learning approach to auctions’, *Journal of economic theory* **82**(1), 65–88.
- Iyer, K., Johari, R. and Sundararajan, M. (2014), ‘Mean field equilibria of dynamic auctions with learning’, *Management Science* **60**(12), 2949–2970.
- Jiang, C., Beck, C. L. and Srikant, R. (2014), ‘Bidding with limited statistical knowledge in online auctions’, *SIGMETRICS Perform. Eval. Rev.* **41**(4), 38–41.
- Karande, C., Mehta, A. and Srikant, R. (2013), Optimizing budget constrained spend in search advertising, *in* ‘Sixth ACM International Conference on Web Search and Data Mining, WSDM 2013, Rome, Italy, February 4-8, 2013’, pp. 697–706.

- Kleinberg, R. (2005), A multiple-choice secretary algorithm with applications to online auctions, in ‘Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms’, SODA ’05, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, pp. 630–631.
- Kushner, H. J. and Yin, G. G. (2003), *Stochastic approximation and recursive algorithms and applications*, Springer-Verlag, New York.
- Lai, T. (2003), ‘Stochastic approximation’, *The Annals of Statistics* **31**, 391–406.
- Martello, S. and Toth, P. (1990), *Knapsack Problems: Algorithms and Computer Implementations*, John Wiley & Sons, Inc., New York, NY, USA.
- Mirroknii, V. S., Gharan, S. O. and Zadimoghaddam, M. (2012), Simultaneous approximations for adversarial and stochastic online budgeted allocation, in ‘Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms’, SODA ’12, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, pp. 1690–1701.
- Nedic, A. and Ozdaglar, A. (2009), ‘Distributed subgradient methods for multi-agent optimization’, *IEEE Transactions on Automatic Control* **54**(1), 48–61.
- Nemirovski, A., Juditsky, A., Lan, G. and Shapiro, A. (2009), ‘Robust stochastic approximation approach to stochastic programming’, *SIAM Journal on Optimization* **19**(4), 1574–1609.
- Nemirovski, A. and Yudin, D. (1983), *Problem Complexity and Method Efficiency in Optimization*, John Wiley, New York.
- Rosen, J. B. (1965), ‘Existence and uniqueness of equilibrium points for concave n-person games’, *Econometrica* **33**(3), 520–534.
- Sani, A., Neu, G. and Lazaric, A. (2014), Exploiting easy data in online optimization, pp. 810–818.
- Seldin, Y. and Slivkins, A. (2014), One practical algorithm for both stochastic and adversarial bandits., pp. 1287–1295.
- Shalev-Shwartz, S. (2012), ‘Online learning and online convex optimization’, *Foundations and Trends in Machine Learning* **4**(2), 107–194.
- Talluri, K. T. and van Ryzin, G. J. (2004), *The Theory and Practice of Revenue Management*, International Series in Operations Research & Management Science, Vol. 68, Springer.
- Talluri, K. and van Ryzin, G. (1998), ‘An analysis of bid-price controls for network revenue management’, *Management Science* **44**(11), 1577–1593.
- Tsybakov, A. B. (2008), *Introduction to Nonparametric Estimation*, Springer.
- Weed, J., Perchet, V. and Rigollet, P. (2016), ‘Online learning in repeated auctions’, *Proceedings of the Conference on Learning Theory (COLT’16)*.
- Yao, A. C. C. (1977), ‘Probabilistic computations: Toward a unified measure of complexity’, *Foundations of Computer Science (FOCS), 18th IEEE Symposium on* pp. 222–227.
- Zhou, Y., Chakrabarty, D. and Lukose, R. (2008), Budget constrained bidding in keyword auctions and online knapsack problems, Springer Berlin Heidelberg, pp. 566–576.
- Zinkevich, M. (2003), ‘Online convex programming and generalized infinitesimal gradient ascent’, *20th International Conference on Machine Learning* pp. 928–936.