Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Contextual Standard Auctions with Budgets: Revenue Equivalence and Efficiency Guarantees

Santiago Balseiro, Christian Kroer, Rachitesh Kumar

https://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–18
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Contextual Standard Auctions with Budgets: Revenue Equivalence and Efficiency Guarantees

Santiago Balseiro,[a,*] Christian Kroer,[a] Rachitesh Kumar[a]

[a] Columbia University, New York, New York 10027

*Corresponding author

Contact: srb2155@columbia.edu, https://orcid.org/0000-0002-0012-3292 (SB); christian.kroer@columbia.edu, https://orcid.org/0000-0002-9009-8683 (CK); rk3068@columbia.edu, https://orcid.org/0000-0001-7803-6861 (RK)

Received: April 14, 2022
Revised: September 14, 2022
Accepted: December 29, 2022
Published Online in Articles in Advance: April 6, 2023

https://doi.org/10.1287/mnsc.2023.4719

Abstract. The internet advertising market is a multibillion dollar industry in which advertisers buy thousands of ad placements every day by repeatedly participating in auctions. An important and ubiquitous feature of these auctions is the presence of campaign budgets, which specify the maximum amount the advertisers are willing to pay over a specified time period. In this paper, we present a new model to study the equilibrium bidding strategies in standard auctions, a large class of auctions that includes first and second price auctions, for advertisers who satisfy budget constraints on average. Our model dispenses with the common yet unrealistic assumption that advertisers' values are independent and instead assumes a contextual model in which advertisers determine their values using a common feature vector. We show the existence of a natural value pacing–based Bayes–Nash equilibrium under very mild assumptions. Furthermore, we prove a revenue equivalence showing that all standard auctions yield the same revenue even in the presence of budget constraints. Leveraging this equivalence, we prove price of anarchy bounds for liquid welfare and structural properties of pacing-based equilibria that hold for all standard auctions. In recent years, the internet advertising market has adopted first price auctions as the preferred paradigm for selling advertising slots. Our work, thus, takes an important step toward understanding the implications of the shift to first price auctions in internet advertising markets by studying how the choice of the selling mechanism impacts revenues, welfare, and advertisers' bidding strategies.

History: Accepted by Itai Ashlagi, revenue management and market analytics.
Supplemental Material: The online appendix is available at https://doi.org/10.1287/mnsc.2023.4719.

Keywords: first price auctions • contextual value models • budget constraints • equilibria in auctions • revenue equivalence • internet advertising

1. Introduction

In 2019, the revenue from selling internet ads in the United States surpassed $129 billion.[1] A large fraction of these are sold on ad platforms operated by tech giants such as Google, Facebook, and Twitter. These platforms facilitate the sale of ads by acting as intermediaries between advertisers and publishers. Millions of ad slots are sold every day using auctions in which advertisers bid based on user-specific information (such as geographical location, cookies, and historical activity, among others). The advertisers repeatedly participate in these auctions with the aim of using their advertising budget to maximize their reach through a combination of user-specific targeting and bid optimization. The presence of budgets introduces significant challenges as it links different auctions together.

With billions of dollars at stake, the auction format plays a crucial role. In recent years, a major shift has occurred toward using first price auctions as the preferred mode of selling display ads as opposed to the earlier standard of using second price auctions. For example, in 2019, Google, which is one of the industry leaders, announced a shift to the first price auction format for its ad exchange.[2] In 2020, Twitter also made the move to first price auctions for the sale of mobile app advertising slots.[3] First price auctions typically lead to more complicated bidding behavior because, unlike second price auctions, truthful bidding is not an equilibrium in the first price setting.

This paper attempts to capture the salient features of these display ad auctions with a focus on the newly adopted first price auctions. Whereas equilibrium behavior in first price auctions is studied extensively, very little attention is given to the effects of budget constraints and user-specific information. Budget constraints span the auctions, which means that advertisers must strategize about their bids across all auctions simultaneously. User-specific information leads to correlation between the valuations that different advertisers have for a particular ad opportunity, whereas the literature largely focuses on independent and identically distributed (i.i.d.) valuations. Our paper aims to shed some light on these aspects by introducing and analyzing a framework for general standard auctions, including first price auctions, that incorporates budgets and context-based valuations. In particular, the main questions we tackle are as follows: How does the auction format affect the equilibrium strategies of budget-constrained bidders? How does the auction format impact the revenue of the ad platform and the efficiency of the market?

1.1. Main Contributions

We incorporate the availability of user-specific information (that is common to all advertisers) via a contextual valuation model, which allows us to capture correlation in values. User information and buyer targeting criteria are modeled as vectors with the value that an advertiser gains from the advertiser's ad being shown to the user being given by the inner product of these vectors or a function thereof. Each advertiser has a (possibly nonbinding) budget that must be satisfied in expectation. Such budget constraints are well-motivated in practice because of the large number of auctions and are the subject of study in previous work on budget management (Gummadi et al. 2011; Abhishek and Hosanagar 2013; Balseiro et al. 2015, 2021). Our main contribution is to introduce a framework that allows for the study of standard auctions, which are auctions in which the highest bidder wins, in the presence of budget constraints and contextual valuations. To the best of our knowledge, this is the first analysis of standard auctions in the presence of average budget constraints.

Typically, the targeting criteria and budget of an advertiser are not known to the advertiser's competitors. This motivates us to model the participation of advertisers in the market as a nonatomic game of incomplete information in which each advertiser assumes that the other advertisers are being drawn from some common underlying distribution. In this game, the decision problem faced by each advertiser is to optimize the advertiser's utility and satisfy the advertiser's budget constraint in expectation. This expectation is taken over all the potential auctions the advertiser could end up participating in; that is, the expectation is over users and competing advertisers. Our nonatomic game allows us to sidestep the possibility of multiple buyers tying in the auction and leads to simple and intuitive equilibrium strategies.

1.1.1. Equilibrium Analysis. A contribution of this paper is to prove the existence of a remarkably simple Bayes–Nash equilibrium (BNE) strategy using a novel topological argument. In our nonatomic model, there is a continuum of advertiser types, and a strategy for each advertiser type is a function that maps contexts to bids. Directly proving existence of an equilibrium in this complicated strategy space in the presence of budget constraints turns out to be difficult. We sidestep this difficulty by establishing strong duality for the constrained nonconvex optimization problem faced by each advertiser type and characterizing the primal optimum in terms of the dual optimum.

We propose a remarkably simple class of strategies, which we dub *value pacing–based* strategies. These strategies build on the symmetric equilibrium strategies of the standard i.i.d. setting, inheriting their interpretability in the process. A value pacing–based strategy recommends that each advertiser should shade the advertiser's value by a multiplicative factor to manage the advertiser's budget and then bid using the symmetric equilibrium strategy from the standard i.i.d. setting—as the advertiser would in the absence of budgets—but assuming that competitors' values are also paced. This naturally extends multiplicative bid pacing/shading, which is one of the several ways budgets are managed in practice, to nontruthful auctions (Balseiro et al. 2021; Conitzer et al. 2022a, b). To the best of our knowledge, our value-pacing approach is the first to show optimal pacing-based strategies outside of truthful auctions.

Our nonatomic game has a pacing (dual) multiplier for each buyer type, which are uncountably many in cardinality. This leads to an infinite-dimensional equilibrium space even after moving to the simpler dual space. In infinite dimensions, establishing even the simple prerequisites of any fixed-point theorem, namely, compactness and continuity, can be an ordeal, one which requires careful topological arguments. Whereas other papers also analyze equilibrium strategies in the dual space (see, e.g., Gummadi et al. 2011, Balseiro et al. 2015), these consider settings with finitely many pacing multipliers in which establishing compactness is a trivial task. The main technical contribution of this paper is twofold: (i) choosing the right topological space for the pacing multipliers based on their monotonicity properties and (ii) establishing compactness and continuity in this carefully chosen space. As we discuss in Section 3.3, this choice of topology is far from obvious. In fact, to the best of our knowledge, all of the topologies used in standard fixed-point arguments for infinite-dimensional spaces (see Aliprantis and Border 2006 for examples) prove insufficient in the setting we consider, which compels us to carefully exploit the structural properties of pacing and work with the topological space of multivariate functions of bounded variation. We believe the tools developed in this paper might be useful in other nonatomic games.

1.1.2. Standard Auctions and Revenue Equivalence. Our framework accommodates anonymous auction formats in which the highest bidder wins, such as second price and all-pay auctions (even in the presence of reserve prices). In its full generality, it acts as a powerful black box: it takes as input any Bayes–Nash equilibrium for the well-studied standard i.i.d. setting, composes it with value pacing, and outputs a Bayes–Nash equilibrium for our model. Surprisingly, we show that, for a fixed distribution over advertisers and users, the same multiplicative factors can be used by the advertisers to shade their values in the equilibrium strategies for all standard auctions. This fact allows us to compare revenues across auction formats. We prove that, in the presence of in-expectation budget constraints, the revenue generated in a value pacing–based equilibrium is the same for all standard auctions. This is in sharp contrast to the case when budget constraints are strict, in which revenue equivalence is known not to hold (Che and Gale 1998). In light of the recent shift from second to first price auctions by many ad platforms, the ability to compare budget management in both first and second price auctions is an especially relevant aspect of our framework. A recent paper of Goke et al. (2021) empirically investigates the revenue impact resulting from this recent switch. Goke et al. (2021) find that, after a brief adjustment period, publishers' revenues under first price auctions returned to the same levels as they were under second price auctions before the change. Because existing revenue equivalence results do not apply to the display-ad industry (because of budget constraints and dependencies in valuations), our theory offers the first principled justification for this empirical finding by establishing revenue equivalence in the presence of contextual values and in-expectation budget constraints.

1.1.3. Price of Anarchy (PoA) and Structural Results. We leverage our revenue equivalence result to establish efficiency guarantees and structural properties that hold for all standard auctions. In particular, we prove a (1/2)-lower bound for the price of anarchy of liquid welfare (a notion of efficiency that incorporates budget constraints) for our value pacing–based equilibria. Our result implies that the liquid welfare of a pacing equilibrium is at most 1/2 of the liquid welfare of the best possible allocation. On the structural front, we study how value pacing–based equilibrium strategies change with buyer type, which consists of a weight vector (representing targeting criteria) and a budget. We show that budget-constrained buyers with identical budgets and colinear weights for features get paced to the same value in equilibrium. This shows that any enhancement in the ad quality without changing its targeting criteria, which corresponds to scaling up the weight vector, is futile in the absence of an increase in budget. Moreover, we also study how advertisers should change their targeting criteria (as represented by their weight vector) to maximize their utility.

1.1.4. Numerical Experiments. To test our model, we run numerical experiments after making appropriate discretizations. The outcomes of these experiments are strikingly close to our theoretical predictions. In particular, despite the discontinuities introduced by discretization, budget violations are small, and moreover, the equilibrium strategies are in strong adherence to the structural properties derived theoretically.

1.2. Related Work

Beyond the works already mentioned, there is a large literature on online auctions. We discuss the existing work that is most closely related to ours. In keeping with previous work on auctions, from now on, we use the terms "buyers" and "items" in place of advertisers and users.

Auctions with budget-constrained buyers are modeled in a variety of ways, most of which are focused on second price auctions. From a technical standpoint, the closest to our work is Balseiro et al. (2015), which considers randomly arriving budget-constrained buyers in a fluid mean field setting. They show equilibrium existence for second price auctions in which buyers use pacing-based strategies. Their model assumes a finite-type space and independence of the value distributions of the buyers, whereas our context-based model allows for correlation between buyer values. Several other works also study repeated second price auctions with budgets under various models that are less related to ours (Gummadi et al. 2011, Balseiro and Gur 2019, Chen et al. 2021a, Balseiro et al. 2021, Ciocan and Iyer 2021, Conitzer et al. 2022a). Beyond second price auctions, Aggarwal et al. (2019) consider affine constraints (which include budget constraints as a special case) in multislot truthful auctions; they show existence of a bid-pacing equilibrium under restrictive assumptions. Babaioff et al. (2021) consider a general model of non–quasi-linear buyers participating in mechanisms that are truthful for quasi-linear buyers. Their model also captures the case of budget constraints as a special case. Moreover, they too prove the existence of a pacing-based equilibrium in their model. None of the aforementioned existing work addresses strategic bidding in nontruthful auctions, such as first price auctions with budget-constrained buyers.

Conitzer et al. (2022b) and Borgs et al. (2007) study pacing in a first price context, but both disregard strategic behavior on behalf of the buyers. This is also the case for a long line of work that models repeated auctions among budget-constrained buyers as an online matching problem with capacity constraints (see Mehta 2013 for a survey).

Another direction of research considers buyers with ex post budget constraints (also called strict budget constraints). There, first price (Kotowski 2020), standard (Che and Gale 1998), and optimal (Pai and Vohra 2014) auctions and auctions with combinatorial constraints (Goel et al. 2015) are studied. In contrast to our revenue equivalence results, Che and Gale (1998) show that, with strict budget constraints, first price auctions yield higher revenue than second price auctions. These models are different from our setting, which only requires budget constraints to hold in expectation at the interim stage. In-expectation budget constraints are more appropriate for modeling repeated ad auctions and yield simpler and more interpretable equilibrium strategies.

Contextual models in which values are based on feature vectors are widely used in the multiarmed bandit literature (for example, see Langford and Zhang 2007, Li et al. 2010) and in pricing (Lobel et al. 2018, Chen et al. 2021b, Golrezaei et al. 2021). Vector-based valuation models are also connected to low-rank models, which have received attention in previous market design work (see, e.g., McMahan et al. 2013, Kroer et al. 2022).

Our work also contributes to the literature on equilibrium analysis for nonatomic games. Because of the presence of a continuum of buyer types in our model, the topological arguments we develop bear resemblance to those used in the study of nonatomic games, such as the one addressed in Schmeidler (1973), though continuity is by assumption in Schmeidler (1973), whereas achieving continuity is at the heart of our proof.

2. Model

We consider the setting in which a seller (i.e., the advertising platform) plans to sell an indivisible item to one of n buyers (i.e., the advertisers) using an auction. We adopt a feature-based valuation model for the buyer. More precisely, the item type is represented using a vector α belonging to the set $A \subset \mathbb{R}^d$, where each component of α can be interpreted as a feature. We also refer to α as the context. Each buyer type is represented using a vector (w, B) belonging to the set $\Theta \subset \mathbb{R}^{d+1}$ of possible buyer types, where the last component B denotes the buyer's budget and the first d components w capture the weights the buyer assigns to each of the d features. The value (maximum willingness to pay) that buyer type (w, B) has for item α is given by the inner product $w^T\alpha$. For simplicity of notation and ease of exposition, we state our results under this linear relationship between values and the features, but our model and results can be extended to accommodate nonlinear response functions (such as the logistic function) that are commonly used in practice (see Online Appendix G for a more detailed discussion). We use $\omega = \max_{(w,B)\in\Theta,\alpha\in A} w^T\alpha$ to denote the maximum value that a buyer can have for an item.

We assume that the context of the item to be auctioned is drawn from some distribution F over the set of possible item types A. Furthermore, the type for every buyer is drawn according to some distribution G over the set of possible buyer types Θ, independently of the other buyers and the choice of the item. Note that, by virtue of our context-based valuation model, the values of the n buyers for the item need not be independent. In line with standard models used for Bayesian analysis of auctions, we assume that both G and F are common knowledge and maintain that the realized type vector (w, B) associated with a buyer is the buyer's private information. Our model allows budgets to be random and correlated with the buyer's weight vector. In addition, we assume that buyers are unaware of the type of their competing buyers; this implies budgets are private.

To fix ideas, we first consider the case of a first price auction with reserve prices and then discuss how our results extend to standard auctions in Section 4. In a first price auction, the seller allocates the item to the highest bidder whenever the bid is above the reserve price, and the winning bidder pays the bid. We assume the seller discloses the item type α to the n buyers before bids are solicited from them. As a result, the bid of a buyer on item α can depend on α. We use $r : A \to \mathbb{R}$ to specify the publicly known context-dependent reserve prices, where $r(\alpha)$ denotes the reserve price on item type α.

The budget of a buyer represents an upper bound on the amount the buyer wants to pay in the auction. We only require that each buyer satisfy the buyer's budget constraint in expectation over the item type and competing buyer types. Similar assumptions are made in the literature (see, e.g., Gummadi et al. 2011; Abhishek and Hosanagar 2013; Balseiro et al. 2015, 2021; Conitzer et al. 2022a). The motivation behind this modeling choice is that budget constraints are often enforced, on average, by advertising platforms. For example, Google Ads allows daily budgets to be exceeded by a factor of two in any given day, but over the course of a month, the total expenditure never exceeds the daily budget times the days in the month.[4] In-expectation budget constraints are also motivated by the fact that, in practice, buyers typically participate in a large number of auctions and many buyers use stationary bidding strategies. Thus, by the law of large numbers, our model can be interpreted as collapsing multiple, repeated auctions in which item types are drawn i.i.d. from F into a single one-shot auction with in-expectation constraints.

2.1. Notation

We use $\mathbb{R}_+$ and $\mathbb{R}_{\geq 0}$ to denote the set of strictly positive and nonnegative real numbers, respectively. We use G_w to denote the marginal distribution of w when $(w, B) \sim G$, that is, $G_w(K) := G(\{(w, B) \in \Theta | w \in K\})$ for all Borel sets $K \subset S$. In a similar vein, we use Θ_w to denote the set of $w \in \mathbb{R}^d$ such that $(w, B) \in \Theta$ for some $B \in \mathbb{R}$. (Here, we abuse notation by using w as both a weight vector variable and a subscript to denote the projection of a buyer type onto the first d dimensions). Unless specified otherwise, $\|\cdot\|$ denotes the Euclidean norm.

2.2. Assumptions

We assume that there exist $U, B_{\min} > 0$ such that the set of possible buyer types Θ is given by $\Theta = (0, U)^d \times (B_{\min}, U)$. In a similar vein, we also assume that the set of possible item types A is a subset of the positive orthant $\mathbb{R}^d_+$. We restrict our attention to $d \geq 2$, which is the regime in which our feature vector-based valuation model yields interesting insights. To completely specify the aforementioned probability spaces, we endow A, Θ, and Θ_w with the Lebesgue $\sigma-$algebra. Moreover, we assume that the distributions F and G have density functions. Note that the distribution G can be any distribution on Θ, including one with probability zero on some regions of Θ. Thus, we can address any buyer distribution so long as it has a density and is supported on a bounded subset of the strictly positive orthant with a positive lower bound on the possible budgets. Similarly, F can capture a wide variety of item distributions. It is worth noting that any distribution that lacks a density can be approximated arbitrarily well by a distribution with a density, thereby extending the reach of our results to arbitrary distributions.

2.3. Equilibrium Concept

Consider the decision problem faced by a buyer type $(w, B) \in \Theta$ if we fix the bidding strategies of all competing buyers on all possible item types: the buyer wishes to bid on the items in a way that maximizes the buyer's expected utility and satisfies the buyer's budget constraint in expectation (in which the expectation is taken over items and competing buyers' types). As is true in the well-studied standard, budget-free, i.i.d. setting (Krishna 2009), the buyer's optimal strategy depends on the strategies used by the other buyer types. In the standard setting, the symmetric Bayes–Nash equilibrium is an appealing solution concept for the game formed by these interdependent decision problems faced by the buyers. We adopt a similar approach and define the symmetric Bayes–Nash equilibrium for our model. A strategy $\beta^* : \Theta \times A \to \mathbb{R}_{\geq 0}$ (a mapping that specifies what each buyer type should bid on every item) is a symmetric first price equilibrium (SFPE) if, almost surely (a.s.) over all buyer types, using β^* is an optimal solution to a buyer type's decision problem when all other buyer types also use it.

Definition 1. A strategy $\beta^* : \Theta \times A \to \mathbb{R}_{\geq 0}$ is called a symmetric first price equilibrium if $\beta^*(w, B, \alpha)$ (as a function of α) is an optimal solution to the following optimization problem almost surely with respect to (w.r.t.) $(w, B) \sim G$:

$$\begin{aligned} \max_{b:A\to\mathbb{R}_{\geq 0}} \; & \mathbb{E}_{\alpha,\{\theta_i\}_{i=1}^{n-1}}[(w^T\alpha - b(\alpha))\mathbb{1}\{b(\alpha) \geq \max(r(\alpha), \{\beta^*(\theta_i, \alpha)\}_i)\}] \\ \text{s.t. } & \mathbb{E}_{\alpha,\{\theta_i\}_{i=1}^{n-1}}[b(\alpha)\mathbb{1}\{b(\alpha) \geq \max(r(\alpha), \{\beta^*(\theta_i, \alpha)\}_i)\}] \leq B. \end{aligned}$$

In the buyer's optimization problem, the buyer wins whenever the buyer's bid $b(\alpha)$ is higher than the reserve price $r(\alpha)$ and all competiting bids $\beta^*(\theta_i, \alpha)$ for $i = 1, \ldots, n-1$. Because of the first price auction payment rule, each bidder pays the bidder's bid whenever the bidder wins. For convenience, in the definition, we are using an infeasible tie-breaking rule that allocates the entire good to every highest bidder. This is inconsequential and can be replaced by any arbitrary tie-breaking rule because we later show (see part (d) of Lemma 6 in the Online appendix) that ties are a zero-probability event under our value pacing–based equilibria.

In our solution concept, it is sufficient that advertisers have Bayesian priors over the maximum competing bid $\max_i\{\beta^*(\theta_i, \alpha)\}$ to determine a best response. This is aligned with practice as many advertising platforms provide bidders with historical bidding landscapes, which advertisers can use to optimize their bidding strategies (Bigler 2021).[5] Additionally, we require that budgets are satisfied in expectation over the contexts and buyer types. Connecting back to our repeated auctions interpretation, one can assume competitors' types to be fixed throughout the horizon, whereas contexts are drawn i.i.d. in each auction. In this case, our solution concept is appropriate if buyers cannot observe the types of competitors and, in turn, employ stationary strategies that do not react to the market dynamics. Such stationary strategies are appealing because they deplete budgets smoothly over time and are simple to implement. Moreover, it is previously established that stationary policies approximate well the performance of dynamic policies in nonstrategic settings when the number of auctions is large and the maximum value of each auction is small relative to the budget (see, e.g., Talluri and Van Ryzin 2006).

When the type of bidder is fixed throughout the horizon, a bidder who employs a dynamic strategy could, in principle, profitably deviate by inferring the competitors' types and using this information to optimally shade the bidder's bids. Implementing such strategies in practice is challenging because many platforms do not disclose the identity of the winner or the bids of competitors in real time (as we discuss, they mostly provide historical information that is aggregated over many auctions). Moreover, when the number of bidders is large and each bidder competes with a random subset of bidders, such deviations can be shown to be not profitable using mean-field techniques (see, e.g., Iyer et al. 2014, Balseiro et al. 2015) in our contextual value model as long as values are independent across time. Therefore, our model can be alternatively interpreted as one in which there is a large population of active buyers and each buyer competes with a random subset of buyers. This assumption is well-motivated in the context of internet advertising markets because the number of advertisers actively bidding is typically large and, because of sophisticated targeting technologies, advertisers often participate only in a fraction of all auctions.

2.4. Ties and the Role of Contexts

Before moving onto the proof of existence of SFPE, we shed some light on the role played by contexts in our model and results. The assumption that the feature vectors α are drawn from a distribution F that has a density is necessary for our results to hold. In fact, if there was only one deterministic context, an SFPE may fail to exist: we provide an example in Online Appendix A. The root cause behind the absence of a well-behaved equilibrium in this example is the tension between the proclivity of budgets to cause ties with positive probability (as we demonstrate in Section 6) and the potential lack of equilibria for first price auctions under value distributions that cause ties with a positive probability. Our example in Online Appendix A does admit a symmetric equilibrium for second price auctions, thereby demonstrating the added complexity of dealing with first price auctions.

Issues of tie breaking have previously come up in a line of related work on pacing-based equilibria in second price auctions (Borgs et al. 2007, Balseiro et al. 2015, Babaioff et al. 2021, Conitzer et al. 2022a), in which they were addressed by methods that are some version of randomly perturbing the value of each buyer and enforcing the budget constraint, on average, over these perturbations. This causes ties to become zero-probability events. It is possible to prove our existence and revenue equivalence results for the case of one deterministic context with value perturbations. However, unlike second price auctions in which bidding truthfully is a dominant strategy, value perturbation is not well-suited for first price auctions because, even in the absence of budgets, the first price auction equilibrium strategy depends on the perturbations. Moreover, our structural results (Propositions 3 and 4) may not hold for arbitrary perturbations and require an unjustifiably strong assumption that carefully coordinates the perturbations across buyer types. That said, if one is willing to ignore ties, our results continue to hold for a single deterministic context, and the reader can safely continue with that setting in mind.

3. Existence of Symmetric First Price Equilibrium

In this section, we study the existence of SFPE and show that this existence is achieved by a compelling solution that is interpretable. We do so in several steps. First, we define a natural parameterized class of value pacing–based strategies. Then, assuming that competitors are using a strategy from this class, we establish strong duality for the optimization problem faced by each buyer type and characterize the primal optimum in terms of the dual optimum. This leads to a substantial simplification of the analysis because it allows us to work in the much simpler dual space. Finally, we establish the existence of a value pacing–based SFPE by a fixed-point argument over the space of dual multipliers.

3.1. Value Pacing–Based Strategies

In this paper, pacing refers to multiplicatively scaling down a quantity.[6] Consider a function $\mu : \Theta \to \mathbb{R}_{\geq 0}$, which we refer to as the *pacing function*. We define the *paced weight vector* of a buyer with type (w, B) to be $w/(1+\mu(w,B))$, which is simply the true weight vector w scaled down by the factor $1/(1+\mu(w,B))$. Similarly, we define the *paced value* of a buyer type (w, B) for item α as $w^T\alpha/(1+\mu(w,B))$. We use pacing to ensure that the budget constraints of all buyer types are satisfied and, at the same time, maintain the best response property at equilibrium. The motivation for using pacing as a budget management strategy becomes clear in the next section, in which we show that the best response of a buyer to other buyers using a value pacing–based strategy is to also use a value pacing–based strategy. Before defining the strategy, we set up some preliminaries.

Consider a pacing function $\mu : \Theta \to \mathbb{R}_{\geq 0}$ and an item $\alpha \in A$. Let λ^μ_α denote the distribution of paced values $w^T\alpha/(1+\mu(w,B))$ for item α when $(w,B) \sim G$. Let H^μ_α denote the distribution of the highest value $Y := \max\{X_1, \ldots, X_{n-1}\}$ among $n-1$ buyers when each $X_i \sim \lambda^\mu_\alpha$ is drawn independently for $i \in \{1, \ldots, n-1\}$. Observe that $H^\mu_\alpha((-\infty, x]) = \lambda^\mu_\alpha((-\infty, x])^{n-1}$ for all $\alpha \in A$ because the random variables are i.i.d.

For a given item $\alpha \in A$, when $x \geq r(\alpha)$, define the following bidding function:

$$\sigma^\mu_\alpha(x) := x - \int_{r(\alpha)}^{x} \frac{H^\mu_\alpha(s)}{H^\mu_\alpha(x)} ds,$$

where we interpret $\sigma^\mu_\alpha(x)$ to be zero if $H^\mu_\alpha(x) = 0$. Moreover, when $x < r(\alpha)$, define $\sigma^\mu_\alpha(x) := x$ (we make this choice to ensure that no value below the reserve price gets mapped to a bid above the reserve price, maintaining continuity). Note that $\sigma^\mu_\alpha(x) = \mathbb{E}[\max(Y, r) | Y < x]$. If λ^μ_α has a density, then σ^μ_α is the same as the single-auction equilibrium strategy for a standard first price auction without budgets when the buyer values are drawn i.i.d. from λ^μ_α and the item has a reserve price of $r(\alpha)$ (see, e.g., section 2.5 of Krishna 2009). Our value pacing–based strategy uses $\sigma^\mu_\alpha(x)$ as a building block by composing it with value pacing.

Definition 2. The value pacing–based strategy $\beta^\mu : \Theta \times A \to \mathbb{R}_{\geq 0}$ for pacing function $\mu : \Theta \to \mathbb{R}_{\geq 0}$ is given by

$$\beta^\mu(w, B, \alpha) := \sigma^\mu_\alpha\left(\frac{w^T\alpha}{1+\mu(w,B)}\right) \qquad \forall (w, B) \in \Theta, \alpha \in A.$$

The bid $\beta^\mu(w, B, \alpha)$ is the amount that a non–budget-constrained buyer with type (w, B) would bid on item α if the buyer acted as if the buyer's paced value was the buyer's true value (this is captured by the use of the paced value as the argument for σ^μ_α) and believed that the rest of the buyers were also acting in this way (this is captured by the use of σ^μ_α). Therefore, our strategy has a simple interpretation: bidders pace their values and then bid as in a first price auction in which competitors' values are also paced. Consequently, under our strategy, bidders are shading their values twice: first when determining their paced values $w^T\alpha/(1+\mu(w,B))$ to account for budget constraints and then again when adopting the bidding function σ_α^μ for the first price auction. The bidding strategy σ_α^μ optimally trades off two effects: on the one hand, bidding too close to their paced values leaves no utility to buyers because they pay their bid in case of winning, and on the other hand, bidding too low decreases payments at the expense of also decreasing the chance of winning.

Observe that value pacing–based strategies greatly reduce the degrees of freedom in the system. Instead of specifying a bidding strategy, which is a function, for each buyer type, we only need to specify a scalar $\mu(w,B)$ for each buyer type. In addition, our dual characterization allows us to optimize over the space of all bidding strategies without imposing any restriction on the class of admissible functions. Having defined value pacing–based strategies, we are now ready to state our main existence result.

Theorem 1. *There exists a pacing function $\mu:\Theta\to\mathbb{R}_{\geq 0}$ such that the value pacing–based strategy $\beta^\mu:\Theta\times A\to\mathbb{R}_{\geq 0}$ is an SFPE.*

Before proceeding with the proof of Theorem 1, we note some of its practical prescriptions: (i) It recommends that buyers should pace their value to manage their budgets. As we later show, the equilibrium pacing functions for first price auctions are identical to the ones for second price auctions. This suggests that pacing-based budget-management techniques developed for second price auctions (such as Balseiro and Gur 2019) can be used for first price auctions to compute the paced valued. (ii) Advertising platforms typically provide bidding landscapes to the buyers that allow them to compute the optimal bid for a given value. Given a context α, if $\mathbb{P}_\alpha^\mu$ represents the equilibrium bidding landscape (distribution of highest competing bids), then we have

$$\sigma_\alpha^\mu(x)\in\arg\max_b (x-b)\mathbb{P}_\alpha^\mu(b).$$

Thus, the paced value can be combined with the landscape to compute the optimal bid $\beta^\mu(w,B,\alpha)$.

We provide the proof of Theorem 1 in the remaining sections. First, in Section 3.2, we show that, if all of the competing buyers are assumed to employ a value pacing–based strategy, then strong duality holds for the budget-constrained utility maximization problem faced by each buyer type. This allows us to drastically simplify the equilibrium strategy space of each buyer type from a function (mapping contexts to bids) to a single scalar (the dual variable $\mu(w,B)$). Next, in Section 3.3, we prove the existence of a value pacing–based equilibrium strategy by proving a fixed-point theorem in the dual space of pacing functions. Despite our simplifying move to the dual space, establishing a fixed point is by no means a straightforward task because we are still left with a dual variable for each buyer type, and there are (uncountable) infinitely many of those. This leads to an infinite-dimensional fixed-point problem that requires careful topological analysis. We find that the commonly employed general-purpose topologies fail for our problem, and this motivates us to carefully exploit the structure of pacing to select the right topology.

3.2. Strong Duality and Best Response Characterization

We start by considering the optimization problem faced by an individual buyer with type (w,B) when all competing buyers use the value pacing–based strategy with pacing function $\mu:\Theta\to\mathbb{R}_{\geq 0}$. Denoting by $Q^\mu(w,B)$ the optimal expected utility of such a buyer, we have

$$\begin{aligned} Q^\mu(w,B)=\max_{b:A\to\mathbb{R}_{\geq 0}}\ &\mathbb{E}_{\alpha,\{\theta_i\}_{i=1}^{n-1}}[(w^T\alpha-b(\alpha))\mathbb{1}\{b(\alpha)\\ &\geq\max(r(\alpha),\{\beta^\mu(\theta_i,\alpha)\}_i)\}]\\ \text{s.t.}\ &\mathbb{E}_{\alpha,\{\theta_i\}_{i=1}^{n-1}}[b(\alpha)\mathbb{1}\{b(\alpha)\\ &\geq\max(r(\alpha),\{\beta^\mu(\theta_i,\alpha)\}_i)\}]\leq B.\end{aligned}$$

Our goal in this section is to show that the value pacing–based strategy put forward in Definition 2 is a best response when competitors are pacing their bids according to a pacing function μ.

Remark 1. Compare $Q^\mu(w,B)$ to the definition of an SFPE (Definition 1) and observe that, if we are able to show that there exists $\mu:\Theta\to\mathbb{R}_{\geq 0}$ such that $\beta^\mu(w,B,\cdot)$ is an optimal solution to $Q^\mu(w,B)$ almost surely w.r.t. $(w,B)\sim G$, then β^μ is an SFPE.

For $\mu:\Theta\to\mathbb{R}_{\geq 0}$ and $(w,B)\in\Theta$, consider the Lagrangian optimization problem of $Q^\mu(w,B)$ in which we move the budget constraint to the objective using the Lagrange multiplier $t\geq 0$. We use t to denote the multiplier of one buyer in isolation to distinguish from μ, which is a function giving a multiplier for every buyer type. Denoting by $q^\mu(w,B,t)$ the dual function, we have that

$$\begin{aligned} q^\mu(w,B,t)&=\max_{b(\cdot)}\mathbb{E}_{\alpha,\{\theta_i\}_{i=1}^{n-1}}[(w^T\alpha-(1+t)b(\alpha))\mathbb{1}\{b(\alpha)\\ &\qquad\geq\max(r(\alpha),\{\beta^\mu(\theta_i,\alpha)\}_i)\}]+tB\\ &=(1+t)\max_{b(\cdot)}\mathbb{E}_{\alpha,\{\theta_i\}_{i=1}^{n-1}}\left[\left(\frac{w^T\alpha}{1+t}-b(\alpha)\right)\mathbb{1}\{b(\alpha)\right.\\ &\qquad\left.\geq\max(r(\alpha),\{\beta^\mu(\theta_i,\alpha)\}_i)\}\right]+tB.\end{aligned}$$

The dual problem of $Q^\mu(w,B)$ is given by $\min_{t\geq 0} q^\mu(w,B,t)$.

The next lemma states that the optimal solution to the Lagrangian optimization problem is a value pacing–based strategy. More specifically, for every pacing function $\mu:\Theta\to\mathbb{R}_{\geq 0}$, buyer type (w,B), and dual multiplier t, the value pacing–based strategy $\sigma_\alpha^\mu(w^T\alpha/(1+t))$ is an optimal solution to the Langrangian relaxation of $Q^{\mu}(w,B)$ corresponding to multiplier t. Note that, in general, t need not be equal to $\mu(w,B)$.

Lemma 1. *For pacing function $\mu:\Theta\to\mathbb{R}_{\geq 0}$, buyer type $(w,B)\in\Theta$, and dual multiplier $t\geq 0$,*

$$\sigma_{\alpha}^{\mu}\left(\frac{w^T\alpha}{1+t}\right)\in\arg\max_{b(\cdot)}\mathbb{E}_{\alpha,\{\theta_i\}_{i=1}^{n-1}}\left[\left(\frac{w^T\alpha}{1+t}-b(\alpha)\right)\mathbb{1}\{b(\alpha)\geq\max(r(\alpha),\{\beta^{\mu}(\theta_i,\alpha)\}_i)\}\right].$$

In the proof of Lemma 1, we actually show something stronger than the statement of Lemma 1: the value pacing–based strategy is optimal pointwise for each α and not just in expectation over α. This follows from the observation that, once we fix an item α, we are solving the best response optimization problem faced by a buyer with value $w^T\alpha/(1+t)$ in the standard i.i.d. setting (Krishna 2009) with competing buyer values being drawn from λ_{α}^{μ} and under the assumption that the competing buyers use the strategy σ_{α}^{μ}. If λ_{α}^{μ} has a strictly positive density, then the optimality of $\sigma_{\alpha}^{\mu}(w^T\alpha/(1+t))$ is a direct consequence of the definition of a symmetric BNE in the standard i.i.d. setting. Even though the standard results cannot be used directly because of the potential absence of a density in the situation outlined, we show that it is possible to adapt the techniques used in the proof of proposition 2.2 of Krishna (2009) to show Lemma 1.

Using Lemma 1, we can simplify the expression for the dual function $q^{\mu}(w,B,t)$. First, note that because σ_{α}^{μ} is nondecreasing, the highest competing bid can be written as

$$\max_{i=1,\ldots,n-1}\{\beta^{\mu}(\theta_i,\alpha)\}=\max_{i=1,\ldots,n-1}\left\{\sigma_{\alpha}^{\mu}\left(\frac{w_i^T\alpha}{1+\mu(\theta_i)}\right)\right\}=\sigma_{\alpha}^{\mu}(Y),$$

where $Y\sim H_{\alpha}^{\mu}$ is the maximum of $n-1$ paced values. Therefore, using that $\sigma_{\alpha}^{\mu}(w^T\alpha/(1+t))$ is an optimal bidding strategy, we get that

$$\begin{aligned}
q^{\mu}(w,B,t)&=(1+t)\mathbb{E}_{\alpha}\mathbb{E}_{Y\sim H_{\alpha}^{\mu}}\left[\left(\frac{w^T\alpha}{1+t}-\sigma_{\alpha}^{\mu}\left(\frac{w^T\alpha}{1+t}\right)\right)\mathbb{1}\left\{\sigma_{\alpha}^{\mu}\left(\frac{w^T\alpha}{1+t}\right)\geq\max(r(\alpha),\sigma_{\alpha}^{\mu}(Y))\right\}\right]+tB\\
&=(1+t)\mathbb{E}_{\alpha}\mathbb{E}_{Y\sim H_{\alpha}^{\mu}}\left[\left(\frac{w^T\alpha}{1+t}-\sigma_{\alpha}^{\mu}\left(\frac{w^T\alpha}{1+t}\right)\right)\mathbb{1}\left\{\frac{w^T\alpha}{1+t}\geq\max(r(\alpha),Y)\right\}\right]+tB\\
&=(1+t)\mathbb{E}_{\alpha}\left[\left(\frac{w^T\alpha}{1+t}-\sigma_{\alpha}^{\mu}\left(\frac{w^T\alpha}{1+t}\right)\right)H_{\alpha}^{\mu}\left(\frac{w^T\alpha}{1+t}\right)\mathbb{1}\left\{\frac{w^T\alpha}{1+t}\geq r(\alpha)\right\}\right]+tB\\
&=(1+t)\mathbb{E}_{\alpha}\left[\mathbb{1}\left\{\frac{w^T\alpha}{1+t}\geq r(\alpha)\right\}\int_{r(\alpha)}^{\frac{w^T\alpha}{1+t}}H_{\alpha}^{\mu}(s)ds\right]+tB,
\end{aligned}$$

where the second equation follows from part (c) of Lemma 6 of the Online appendix, the third from taking expectations with respect to Y, and the last from our formula for σ_{α}^{μ}.

We now present the main result of this section, which characterizes the optimal solution of $Q^{\mu}(w,B)$ in terms of the optimal solution of the dual problem. The idea of using value pacing–based strategies as candidates for the equilibrium strategy owes its motivation to Proposition 1. It establishes that, if all the other buyers are using a value pacing–based strategy with some pacing function $\mu:\Theta\to\mathbb{R}_{\geq 0}$, then a value pacing–based strategy is a best response for a given buyer (w,B).

Proposition 1. *There exists $\Theta'\subset\Theta$ such that $G(\Theta')=1$, and for all pacing functions $\mu:\Theta\to\mathbb{R}_{\geq 0}$ and buyer types $(w,B)\in\Theta'$, if t^* is an optimal solution to the dual problem, that is, if $t^*\in\arg\min_{t^*\geq 0}q^{\mu}(w,B,t)$, then $\sigma_{\alpha}^{\mu}(w^T\alpha/(1+t^*))$ is an optimal solution for the optimization problem $Q^{\mu}(w,B)$.*

In Proposition 1, the pacing parameter t^* used for pacing in the best response can, in general, be different from $\mu(w,B)$. This caveat requires a fixed-point argument to resolve, which is the subject matter of the next section.

Remark 2. Restricting to the measure-one set Θ' is without loss. Recall that, according to Definition 1, a strategy constitutes an SFPE if, almost surely over $(w,B)\sim G$, using β^* is an optimal solution to the optimization problem when all other buyer types also use it. As a consequence of this definition, we show that it suffices to show strong duality for a subset of buyer types $\Theta'\subset\Theta$ such that $G(\Theta')=1$. In the absence of reserve prices $r(\alpha)$ for the items, Proposition 1 holds for all $(w,B)\in\Theta$. Reserve prices introduce some discontinuities in the utility and payment term. The subset $\Theta'\subset\Theta$ captures a collection of buyer types for which these discontinuities are inconsequential and maintains $G(\Theta')=1$.

Observe that $Q^{\mu}(w,B)$ is not a convex optimization problem, so in order to prove the theorem, we cannot appeal to the well-known strong duality results established for convex optimization. Instead, we use theorem 5.1.5 of Bertsekas et al. (1998), which states that, to prove optimality of $\sigma_{\alpha}^{\mu}(w^T\alpha/(1+t^*))$ for $Q^{\mu}(w,B)$, it suffices to show primal feasibility of $\sigma_{\alpha}^{\mu}(w^T\alpha/(1+t^*))$, dual feasibility of t^*, Lagrange optimality of $\sigma_{\alpha}^{\mu}(w^T\alpha/(1+t^*))$ for multiplier t^*, and complementary slackness. Our approach is to show these required properties by combining the differentiability of the dual function with first order optimality conditions for one-dimensional optimization problems. The key observation here is that the derivative of the dual function is equal to the difference between the budget of the buyer and the buyer's expected expenditure. Therefore, at optimality, the first order conditions of the dual problem imply feasibility of the value-based pacing strategy. To prove differentiability, we leverage that, in our game, the distribution of competing bids is absolutely continuous, which is critical for our results to hold.

For $t^* \in \arg\min_{t\geq 0} q^\mu(w,B,t)$, if we apply the first order optimality conditions for an optimization problem with a differentiable objective function over the domain $[0,\infty)$, we get

$$\frac{\partial q^\mu(w,B,t^*)}{\partial t} \geq 0, \qquad t^* \geq 0, \qquad t^* \cdot \frac{\partial q^\mu(w,B,t^*)}{\partial t} = 0.$$

The first condition can be shown to imply primal feasibility, the second implies dual feasibility, and the third implies complementary slackness. Combining this with Lemma 1, which establishes Lagrange optimality, and applying theorem 5.1.5 of Bertsekas et al. (1998) yields Proposition 1. The complete proof of Proposition 1 can be found in Online Appendix B.

3.3. Fixed-Point Argument

In light of Proposition 1, the proof of Theorem 1 (the existence of a value pacing–based SFPE) boils down to showing that there exists a pacing function $\mu:\Theta\to\mathbb{R}_{\geq 0}$ such that almost surely w.r.t. $(w,B)\sim G$, $\mu(w,B)$ is an optimal solution to the dual optimization problem $\min_{t\geq 0} q^\mu(w,B,t)$. In other words, given that everybody else acts according to μ, a buyer (w,B) that wishes to minimize the dual function is best off acting according to μ. More specifically, in Proposition 1, we show that, starting from a pacing function $\mu:\Theta\to\mathbb{R}_{\geq 0}$, if $\mu^*(w,B)$ constitutes an optimal solution to the dual problem $\min_{t\geq 0} q^\mu(w,B,t)$ almost surely w.r.t. $(w,B)\sim G$, then $\sigma_\alpha^\mu(w^T\alpha/(1+\mu^*(w,B)))$ is an optimal solution for the optimization problem $Q^\mu(w,B)$ almost surely w.r.t. $(w,B)\sim G$. In other words, bidding according to σ_α^μ and pacing according to $\mu^*:\Theta\to\mathbb{R}_{\geq 0}$ is a utility-maximizing strategy for buyer $(w,B)\sim G$ almost surely given that other buyers bid according to σ_α^μ with paced values obtained from μ. The following theorem establishes the existence of a pacing function $\mu:\Theta\to\mathbb{R}_{\geq 0}$ for which μ itself fills the role of μ^* in the previous statement, thereby implying the optimality of $\sigma_\alpha^\mu(w^T\alpha/(1+\mu(w,B)))$ almost surely w.r.t. $(w,B)\sim G$.

Proposition 2. *There exists $\mu:\Theta\to\mathbb{R}_{\geq 0}$ such that $\mu(w,B)\in\arg\min_{t\geq 0} q^\mu(w,B,t)$ almost surely w.r.t. $(w,B)\sim G$.*

We prove this statement using an infinite-dimensional fixed-point argument on the space of pacing functions with a carefully chosen topology. Informally, we need to show that the correspondence that maps a pacing function $\mu:\Theta\to\mathbb{R}_{\geq 0}$ to the set of dual-optimal pacing functions $\mu^*:\Theta\to\mathbb{R}_{\geq 0}$ that satisfy $\mu^*(w,B)\in\arg\min_{t\geq 0} q^\mu(w,B,t)$ has a fixed point. However, unlike finite-dimensional fixed-point arguments, establishing the sufficient conditions of convexity and compactness needed to apply infinite-dimensional fixed-point theorems requires a careful topological argument.

Lemma 8 in the online appendix shows that all dual optimal functions $\mu^*:\Theta\to\mathbb{R}_{\geq 0}$ map to a range that is a subset of $[0,\omega/B_{\min}]$. Therefore, any pacing function $\mu:\Theta\to\mathbb{R}_{\geq 0}$ that is a fixed point, that is, satisfies $\mu(w,B)\in\arg\min_{t\geq 0} q^\mu(w,B,t)$ almost surely w.r.t. $(w,B)\sim G$ must also satisfy range$(\mu)\subset[0,\omega/B_{\min}]$. Hence, it suffices to restrict our attention to pacing functions of the form $\mu:\Theta\to[0,\omega/B_{\min}]$.

Consider the set of all potential pacing functions

$$\mathcal{X} = \{\mu\in L_1(\Theta)\,|\,\mu(w,B)\in[0,\omega/B_{\min}]\ \ \forall\,(w,B)\in\Theta\},$$

where $L_1(\Theta)$ is the space of functions $f:\Theta\to\mathbb{R}$ with finite L_1 norm w.r.t. the Lebesgue measure. Here, by the L_1 norm of f w.r.t. the Lebesgue measure, we mean $\|f\|_{L_1}=\int_\Theta |f(\theta)|d\theta$. Our goal is to find a $\mu\in\mathcal{X}$ such that almost surely w.r.t. $(w,B)\sim G$ we have

$$\mu(w,B)\in \underset{t\in[0,\omega/B_{\min}]}{\arg\min}\ q^\mu(w,B,t).$$

Dealing with infinitely many individual optimization problems $\min_{t\in[0,\omega/B_{\min}]} q^\mu(w,B,t)$, one for each $(w,\ B)$, makes the analysis hard. To remedy this issue, we combine these optimization problems by defining the objective $f:\mathcal{X}\times\mathcal{X}\to\mathbb{R}$ for all $\mu,\hat\mu\in\mathcal{X}$ as follows:

$$f(\mu,\hat\mu) := \mathbb{E}_{(w,B)}[q^\mu(w,B,\hat\mu(w,B))].$$

For a fixed $\mu\in\mathcal{X}$, we then get a single optimization problem $\min_{\hat\mu\in\mathcal{X}} f(\mu,\hat\mu)$ over functions in $\mathcal{X}$ instead of one optimization problem for each of the infinitely many buyer types $(w,B)\in\Theta$. Later, in Lemma 5, we show that any optimal solution to the combined optimization problem is also an optimal solution to the individual optimization problems almost surely w.r.t. $(w,B)\sim G$. Thus, shifting our attention to the combined optimization problem is without any loss (because suboptimality on zero-measure sets is tolerable).

With f as before, we proceed to define the correspondence that is used in our fixed-point argument. The optimal solution correspondence $C^*:\mathcal{X}\rightrightarrows\mathcal{X}$ is given by $C^*(\mu):=\mathrm{argmin}_{\hat\mu\in\mathcal{X}} f(\mu,\hat\mu)$ (which could be empty) for all $\mu\in\mathcal{X}$. In Lemma 5, we show that the proof of Proposition 2 boils down to showing that C^* has a fixed point, which is our next step.

Our proof culminates with an application of the Kakutani–Glicksberg–Fan theorem on a suitable version of C^* to show the existence of a fixed point. An application of this result (or any other infinite-dimensional fixed-point theorem) requires intricate topological considerations. In particular, we need to endow $\mathcal{X}$ with a topology that satisfies the following conditions:

I. The set $\mathcal{X}$ is compact and convex and $C^*(\mu)$ is a nonempty subset of $\mathcal{X}$ for all $\mu\in\mathcal{X}$.

II. The correspondence C^* is a Kakutani map; that is, it is upper hemicontinuous, and $C^*(\mu)$ is compact and convex for all $\mu \in \mathcal{X}$.

In the case of infinite dimensions, bounded sets in many spaces, such as the $L_p(\Omega)$ spaces, are not compact. In particular, $\mathcal{X}$ is not compact as a subset of $L_p(\Omega)$ for any $1 \le p \le \infty$. One possible way around it is to consider the weak* topology on $\mathcal{X} \subset L_\infty(\Omega)$ in which bounded sets are compact. This choice runs into trouble because it is difficult to show the upper hemicontinuity of C^* (property II) under the weak convergence notion of the weak* topology. Alternatively, one could impose structural properties and restrict to a subset of $\mathcal{X}$, such as the space of Lipschitz functions, in which both compactness and continuity can be established. The issue with this approach is that the correspondence operator may, in general, not preserve these properties; that is, property I might not hold. For example, even if μ is Lipschitz, $C^*(\mu)$ might not contain any Lipschitz functions.

We strike a delicate balance between properties I and II by picking a space in which we can establish compactness of $\mathcal{X}$ and upper hemicontinuity of C^* and, at the same time, ensuring that $C^*(\mu)$ contains at least one element from this space. It turns out that the right space that works for our proof is the space of bounded variation. To motivate this topology on the space of pacing functions, we state some properties of the "smallest" dual optimal pacing function. For $\mu : \Theta \to [0, \omega/B_{\min}]$, we define $\ell^\mu : \Theta \to [0, \omega/B_{\min}]$ as

$$\ell^\mu(w, B) := \min\left\{ s \in \underset{t \in [0, \omega/B_{\min}]}{\arg\min}\ q^\mu(w, B, t) \right\}$$

for all $(w, B) \in \Theta$. The minimum always exists because $q^\mu(w, B, t)$ is continuous as a function of t (see Corollary 1 in the online appendix for a proof), and the feasible set of the dual problem is compact.

We first show that ℓ^μ varies nicely with w and B along individual components.

Lemma 2. *For $\mu : \Theta \to [0, \omega/B_{\min}]$, the following statements hold:*

1. *$\ell^\mu : \Theta \to [0, \omega/B_{\min}]$ is nondecreasing in each component of w.*
2. *$\ell^\mu : \Theta \to [0, \omega/B_{\min}]$ is nonincreasing as a function of B.*

The proof applies results from comparative statics, which characterize the way the optimal solutions behave as a function of the parameters, to the family of optimization problems $\min_{t \in [0, \omega/B_{\min}]} q^\mu(w, B, t)$ parameterized by $(w, B) \in \Theta$.

Now, we wish to show bounded variation of ℓ^μ. It is a well-known fact that monotonic functions of one variable have finite total variation. Moreover, functions of bounded total variation also form the dual space of the space of continuous functions with the L_∞ norm, which allows us to invoke the Banach–Alaoglu theorem to establish compactness in the weak* topology. These results for single variable functions, although not directly applicable to the multivariable setting, act as a guide in choosing the appropriate topology for our setting.

Because pacing functions take as input several variables, we need to look at multivariable generalizations of total variation. To this end, we state one of the standard definitions (there are multiple equivalent ones) of total variation for functions of several variables (see section 5.1 of Evans and Gariepy 2015) and then follow it up by a lemma that gives a bound on the total variation of the component-wise monotonic function ℓ^μ.

Definition 3. For an open subset $\Omega \subset \mathbb{R}^n$, the total variation of a function $u \in L_1(\Omega)$ is given by

$$V(u, \Omega) := \sup\left\{ \int_\Omega u(\omega)\, \mathrm{div}\ \phi(\omega) d\omega \,\middle|\, \phi \in C_c^1(\Omega, \mathbb{R}^n), \|\phi\|_\infty \le 1 \right\},$$

where $C_c^1(\Omega, \mathbb{R}^n)$ is the space of continuously differentiable vector functions ϕ of compact support contained in Ω and $\mathrm{div}\ \phi = \sum_{i=1}^n \frac{\partial \phi_i}{\partial x_i}$ is the divergence of ϕ.

Lemma 3. *For any pacing function $\mu : \Theta \to [0, \omega/B_{\min}]$, the following statements hold:*

1. *$\ell^\mu \in L_1(\Theta)$.*
2. *$V(\ell^\mu, \Theta) \le V_0$, where $V_0 := (d+1)U^{d+1}\omega/B_{\min}$ is a fixed constant.*

Motivated by this lemma, we define the set of pacing functions that allow us to use our fixed-point argument. Define $\mathcal{X}_0 = \{\mu \in \mathcal{X} \mid V(\mu, \Theta) \le V_0\}$ to be the subset of pacing functions with variation at most V_0. Note that $\ell^\mu \in \mathcal{X}_0$. Define $C_0^* : \mathcal{X}_0 \rightrightarrows \mathcal{X}_0$ as $C_0^*(\mu) := \arg\min_{\hat{\mu} \in \mathcal{X}_0} f(\mu, \hat{\mu})$ for all $\mu \in \mathcal{X}_0$. We now state the properties satisfied by $\mathcal{X}_0$ that make it compatible with the Kakutani–Fan–Glicksberg fixed-point theorem.

Lemma 4. *The following statements hold:*

1. *$\mathcal{X}_0$ is nonempty, compact, and convex as a subset of $L_1(\Theta)$.*
2. *$f : \mathcal{X}_0 \times \mathcal{X}_0 \to \mathbb{R}$ is continuous when $\mathcal{X}_0 \times \mathcal{X}_0$ is endowed with the product topology.*
3. *$C_0^* : \mathcal{X}_0 \rightrightarrows \mathcal{X}_0$ is upper hemicontinuous with nonempty, convex, and compact values.*

Finally, with this lemma in place, we can apply the Kakutani–Fan–Glicksberg theorem to establish the existence of a $\mu \in \mathcal{X}_0$ such that $\mu \in C_0^*(\mathcal{X}_0)$. The following lemma completes the proof of Proposition 2 by showing that the fixed point is also almost surely optimal for each type. It follows from the fact that, for $\mu \in \mathcal{X}_0$ that satisfy $\mu \in C_0^*(\mu)$, we have $\ell^\mu \in C_0^*(\mu)$.

Lemma 5. *If $\mu \in C_0^*(\mu) = \arg\min_{\hat{\mu} \in \mathcal{X}_0} f(\mu, \hat{\mu})$, then $\mu(w, B)$ is almost surely optimal for each type; that is, $\mu(w, B) \in \arg\min_{t \in [0, \omega/B_{\min}]} q^\mu(w, B, t)$ a.s. w.r.t. $(w, B) \sim G$.*

As mentioned earlier, Proposition 2 combined with Proposition 1 implies Theorem 1.

4. Standard Auctions and Revenue Equivalence

In this section, we move beyond first price auctions and generalize our results to anonymous standard auctions with reserve prices. An auction $\mathcal{A} = (Q, M)$ with allocation rule $Q : \mathbb{R}^n_{\geq 0} \to [0,1]^n$, payment rule $M : \mathbb{R}^n_{\geq 0} \to \mathbb{R}^n_{\geq 0}$, and reserve price r is called an anonymous standard auction if the following conditions are satisfied:

• Highest bidder wins: When the buyers bid $(b_1, \ldots, b_n)$, the allocation received by buyer i is given by $Q_i(b_1, \ldots, b_n) = \mathbb{1}(b_i \geq r, b_i \geq b_j \ \forall j \in [n])$ for all $i \in [n]$.

• Anonymity: The payments made by a buyer do not depend on the identity of the buyer. More formally, if the buyers bid $(b_1, \ldots, b_n)$, then for any permutation π of $[n]$ and buyer $i \in [n]$, we have $M_i(b_1, \ldots, b_n) = M_{\pi(i)}(b_{\pi(1)}, \ldots, b_{\pi(n)})$; that is, the payment made by the ith buyer before the bids are permuted equals the payment made by the bidder $\pi(i)$ after the bids have been permuted.

As in our definition of SFPE, we are using an infeasible tie-breaking rule that allocates the entire good to every highest bidder. As with SFPE, ties are a zero-probability event under our value pacing–based equilibria, and our results hold for arbitrary tie-breaking rules.

For consistency of notation, we modify the preceding notation slightly to better match the one used in previous sections. Exploiting the anonymity of auction $\mathcal{A}$, we denote the payment made by a buyer who bids b when the other $n-1$ buyers bid $\{b_i\}_{i=1}^{n-1}$ by $M(b, \{b_i\}_{i=1}^{n-1})$; that is, we use the first argument for the bid of the buyer under consideration and the other arguments for the competitors' bids. Also, as the reserve price completely determines the allocation rule of a standard auction, in the rest of the section, we omit the allocation rule when discussing anonymous standard auctions and represent them as a tuple $\mathcal{A} = (r, M)$ of reserve price and payment rule.

To avoid delving into the inner workings of the auction, we assume the existence of an oracle that takes as an input an atomless distribution $\mathcal{H}$ over $[0, \omega]$ and outputs a bidding strategy $\psi^{\mathcal{H}} : [0, \omega] \to \mathbb{R}$, satisfying the following properties:

1. The strategy $\psi^{\mathcal{H}}$ is a single-auction equilibrium for the auction $\mathcal{A}$ when the values are drawn i.i.d. from $\mathcal{H}$, that is, $\psi^{\mathcal{H}}(x) \in \arg\max_{b \geq 0} \mathbb{E}_{X_i \sim \mathcal{H}}[x \mathbb{1}\{b \geq \max(r, \{\psi^{\mathcal{H}}(X_i)\}_i)\} - M(b, \{\psi^{\mathcal{H}}(X_i)\}_i)]$.
2. The strategy $\psi^{\mathcal{H}}(x)$ is nondecreasing in x, and $\psi^{\mathcal{H}}(x) \geq r$ if and only if $x \geq r$.
3. The payoff for a bidder who has zero value for the object is zero at the single-auction equilibrium.
4. The distribution of $\psi^{\mathcal{H}}(x)$ when $x \sim \mathcal{H}$ is atomless.

Our results produce a pacing-based equilibrium bidding strategy for budget-constrained buyers by invoking $\psi^{\mathcal{H}}$ as a black box. To make the discussion more concrete, let $\mathcal{A}$ be a second price auction with reserve price r. For a given atomless distribution $\mathcal{H}$, define $\psi^{\mathcal{H}}(v) = v$ to be the truthful bidding strategy. Then, $\psi^{\mathcal{H}}$ is a single-auction equilibrium because bidding truthfully is a dominant strategy in second price auctions. Moreover, $\psi^{\mathcal{H}}$ is nondecreasing, $\psi^{\mathcal{H}}(x) \geq r$ if and only if $x \geq r$, a bidder with zero value bids zero to attain a payoff of zero, and finally the distribution of $\psi^{\mathcal{H}}(x)$ when $x \sim \mathcal{H}$ is simply $\mathcal{H}$, which is atomless. Thus, second price auctions with reserve prices satisfy the assumptions.

In our analysis, we allow the seller to condition on the feature vector and choose a different mechanism for each context $\alpha \in A$. Let $\{\mathcal{A}_\alpha = (r(\alpha), M_\alpha)\}_{\alpha \in A}$ be a family of anonymous standard auctions such that $\alpha \longmapsto r(\alpha)$ is measurable. Moreover, suppose that, for any measurable bidding function $\alpha \longmapsto b(\alpha)$ and any collection of measurable competing bidding functions $\alpha \longmapsto b_i(\alpha)$ for $i \in [n-1]$, the payment function $\alpha \longmapsto M_\alpha(b(\alpha), \{b_i(\alpha)\}_{i=1}^{n-1})$ is also measurable. We define the equilibrium notion for the family $\{\mathcal{A}_\alpha\}_{\alpha \in A}$ of anonymous standard auctions.

Definition 4. A strategy $\beta^* : \Theta \times A \to \mathbb{R}$ is called a symmetric equilibrium for the family of standard auctions $\{\mathcal{A}_\alpha\}_{\alpha \in A}$ if $\beta^*(w, B, \alpha)$ (as a function of α) is an optimal solution to the following optimization problem almost surely w.r.t. $(w, B) \sim G$:

$$\begin{aligned} \max_{b : A \to \mathbb{R}_{\geq 0}} \quad & \mathbb{E}_{\alpha, \{\theta_i\}_{i=1}^{n-1}}[w^T \alpha \mathbb{1}\{b(\alpha) \geq \max(r(\alpha), \{\beta^*(\theta_i, \alpha)\}_i)\} \\ & - M_\alpha(b(\alpha), \{\beta^*(w_i, B_i, \alpha)\}_i)] \\ \text{s.t.} \quad & \mathbb{E}_{\alpha, \{\theta_i\}_{i=1}^{n-1}}[M_\alpha(b(\alpha), \{\beta^*(w_i, B_i, \alpha)\}_i)] \leq B. \end{aligned}$$

Observe that this definition reduces to Definition 1 if we take $\{\mathcal{A}_\alpha\}_{\alpha \in A}$ to be the set of first price auctions with reserve price $r(\alpha)$. Next, we show that the equilibrium existence and characterization results of the previous sections apply to all standard auctions that satisfy the required assumptions. To do this, we first need to define value-pacing strategies for anonymous standard auctions. These are a natural generalization of the value pacing–based strategies used for first price auctions.

Recall that, for a pacing function, $\mu : \Theta \to \mathbb{R}_{\geq 0}$ and $\alpha \in A$, λ^μ_α denotes the distribution of paced values for item α, and H^μ_α denotes the distribution of the highest value for α among $n-1$ buyers. For ease of notation, we use ψ^μ_α to denote the single-auction equilibrium strategy for auction $\mathcal{A}_\alpha$ when values are drawn from $\mathcal{H} = \lambda^\mu_\alpha$ or, more formally, $\psi^\mu_\alpha := \psi^{\lambda^\mu_\alpha}$. For a pacing function $\mu : \Theta \to \mathbb{R}_{\geq 0}$, $(w, B) \in \Theta$ and $\alpha \in A$, define

$$\Psi^\mu(w, B, \alpha) := \psi^\mu_\alpha\left(\frac{w^T \alpha}{1 + \mu(w, B)}\right), \qquad (1)$$

to be our candidate equilibrium strategy. This strategy is well-defined because, by Lemma 6 of the Online appendix, λ^μ_α is atomless almost surely w.r.t. α. As before, the bid $\Psi^\mu(w, B, \alpha)$ is the amount a non–budget-constrained buyer with type (w, B) bids on item α if the buyer's paced value is the buyer's true value when competitors are pacing their values accordingly. In other words, bidders in the proposed equilibrium first pace their values and then bid according to the single-auction equilibrium of auction $\mathcal{A}_\alpha$ in which competitors' values are also paced.

With the definition of value pacing–based strategies in place, we can now state the main result of this section. Recall that, $C_0^*: \mathcal{X}_0 \rightrightarrows \mathcal{X}_0$ is given by $C_0^*(\mu) := \arg\min_{\hat{\mu}\in\mathcal{X}_0} f(\mu, \hat{\mu})$ for all $\mu \in \mathcal{X}_0$, where f is the expected dual function in the case of a first price auction as defined in Section 3.3.

Theorem 2 (Revenue and Pacing Equivalence). *For any pacing function $\mu \in \mathcal{X}_0$ such that $\mu \in C_0^*(\mu)$ is an equilibrium pacing function for first price auctions, the value pacing–based strategy $\Psi^\mu : \Theta \times A \to \mathbb{R}_{\geq 0}$ is a symmetric equilibrium for the family of auctions $\{\mathcal{A}_\alpha\}_{\alpha\in A}$. Moreover, the expected payment made by buyer θ under this equilibrium strategy is equal to the expected payment made by buyer θ in first price auctions under the equilibrium strategy $\beta^\mu : \Theta \times A \to \mathbb{R}_{\geq 0}$, that is,*

$$\mathbb{E}_{\alpha,\{(\theta_i)\}_{i=1}^{n-1}}[M_\alpha(\Psi^\mu(\theta,\alpha), \{\Psi^\mu(\theta_i,\alpha)\}_i)]$$
$$= \mathbb{E}_{\alpha,\{(\theta_i)\}_{i=1}^{n-1}}[\beta^\mu(\theta,\alpha)\,\mathbb{1}\{\beta^\mu(\theta,\alpha) \geq \max(r(\alpha), \{\beta^\mu(\theta_i,\alpha)\}_i)\}].$$

The key step in the proof involves showing that the dual of the budget-constrained utility-optimization problem faced by a buyer is identical for all standard auctions when the other buyers use the equilibrium strategy Ψ^μ of the standard auction under consideration. To establish this key step, we exploit the separable structure of the Lagrangian optimization problem and apply the known utility equivalence result for standard auctions in the single-auction i.i.d. setting once for each item $\alpha \in A$. Then, we establish the analogue of Proposition 1 for standard auctions. Combining this with $\mu \in C_0^*(\mu)$ yields Theorem 2.

Our revenue equivalence relies on three critical assumptions: risk neutrality, independence of weight vectors, and symmetry. As in the classic setting, revenue equivalence fails if buyers are risk-averse (see, e.g., Krishna 2009). We emphasize that, in contrast to the classic revenue equivalence result, buyers' values $w^T\alpha$ are not independent. Our result does require that weight vectors are independent across buyers. Buyers in our model are ex ante homogeneous because buyer types are drawn from the same population. We remark, however, that buyers are heterogenous in the interim sense: the buyers competing in an auction can have different budgets and weight vectors. Revenue equivalence fails if buyers are ex ante heterogenous, that is, if competitors are drawn from different populations.

Before ending this section, we state some important implications of Theorem 2. If the pacing function μ allows buyers to satisfy their budget constraints in some standard auction, then the same pacing function μ allows buyers to satisfy their budgets in every other standard auction. In other words, the equilibrium pacing functions are the same for all standard auctions. This means that, in order to calculate an equilibrium pacing function μ that satisfies $\mu \in C_0^*(\mu)$, it suffices to compute it for any standard auction (in particular, one could consider a second price auction for which bidding truthfully is a dominant-strategy equilibrium in the absence of budget constraints). This fact is especially pertinent in view of the recent shift in auction format used for selling display ads from second to first price auctions because it states that, in equilibrium, buyers can use the same pacing function even after the change. Moreover, the same pacing function continues to work even if the family $\{\mathcal{A}_\alpha\}_{\alpha\in A}$ is an arbitrary collection of first and second price auctions (or any other combination of standard auctions); that is, Theorem 2 states that, not only can one pacing function be used to manage budgets in first and second price auctions, the same pacing function also works in the intermediate transition stages in which buyers may potentially participate in some mixture of these auctions.

Another important takeaway is that all standard auctions with the same allocation rule yield the same revenue to the seller. We remark, however, that the revenue of the seller does depend on the allocation, and the seller could, thus, maximize revenue by optimizing over the reserve prices. We leave the question of optimizing the auction design as a future research direction.

The revenue equivalence in the presence of in-expectation budget constraints is driven by the invariance of the pacing function over all standard auctions and the classic revenue equivalence result for the unconstrained i.i.d. setting, which shows that—on average—payments are the same across standard auctions. Whereas revenue equivalence is known to hold for standard auctions without budget constraints, Che and Gale (1998) show that, when budget constraints are hard, first price auctions lead to higher revenue than second price auctions. The intuition for their result is that, because bids are higher in second than in first price auctions, hard budget constraints are more likely to bind in the former, which reduces the seller's revenue. Surprisingly, Theorem 2 shows that, when budget constraints are in expectation (and values are feature-based), we recover revenue equivalence. To better understand the difference between the two types of constraints, consider the following example.

Example 1. Consider two buyers with values drawn uniformly from the unit interval $[0, 1]$. Moreover, let the budget of the buyer with value v be given by $0.5 + \epsilon v$ for some small $\epsilon > 0$. First, observe that, in the absence of budget constraints, bidding truthfully is a dominant strategy in a second price auction, and bidding half of one's value is a Bayes–Nash equilibrium in a first price auction. Moreover, from the standard revenue-equivalence result, a buyer with value x spends $x^2/2$ in expectation over the other buyer's type in both auctions. Now, because this expected expenditure is less than $1/2$ for all types, the in-expectation budget constraints are nonbinding and the equilibria remain unchanged even when in-expectation budget constraints are imposed. On the other hand, consider the case when the budget constraints are hard. The first price auction equilibrium remains unchanged because every buyer type bids less than 0.5, so the constraint is always satisfied. But, for second price auctions, this is not the case: with hard budget constraints, the equilibrium strategy for the buyers is to bid the minimum of their value and budget, thereby leading to lower revenue compared with the truthful-bidding equilibrium.

We conclude this section with a discussion of extensions and alternative models. First, even though we only consider anonymous standard auctions in this work, our equilibrium existence and revenue equivalence results can be extended to other anonymous allocation rules Q, which (i) admit an oracle that outputs an equilibrium bidding strategy for traditional i.i.d. settings and satisfies properties 1–4 listed at the beginning of this section and (ii) lead to continuous nondecreasing interim-allocation rules for every buyer–item pair when other buyers follow a value pacing–based strategy analogous to the one defined in Equation (1). Second, the argument developed in the section also implies the existence of value pacing–based equilibria and revenue equivalence for standard auctions in the symmetric special case of the models studied in Balseiro et al. (2015, 2021), which consider buyers with ex ante budget constraints that hold in expectation over a buyer's own value and the values of others (see Online Appendix C.1 for a detailed description).

5. Worst Case Efficiency Guarantees

In this section, we use our framework to characterize the price of anarchy, that is, the worst case ratio of the efficiency of a pacing equilibrium relative to the efficiency of the best possible allocation. We measure efficiency of an allocation using the notion of liquid welfare introduced by Dobzinski and Leme (2014), which captures the maximum revenue that can be extracted by a seller who knows the values in advance. We use liquid welfare as a measure of efficiency instead of social welfare because the latter can have arbitrarily small price of anarchy (see Online Appendix D for an example). Throughout this section, we assume that the reserve price is zero for each item; that is, $r(\alpha) = 0$ for all $\alpha \in A$.

We begin by defining the appropriate notion of liquid welfare of an allocation for our model motivated by the original definition of Dobzinski and Leme (2014). Here, an allocation simply refers to a measurable function $x : A \times \Theta^n \to \Delta^n$, where $\Delta^n = \{y \in \mathbb{R}^n_+ \mid \sum_{k=1}^n y_k = 1\}$ is the n-simplex and $x_i(\alpha, \vec{\theta})$ denotes the fraction of the item α allocated to buyer i when the buyer types are given by the profile $\vec{\theta} = (\theta_1, \ldots, \theta_n)$. In our setting, the liquid welfare of a buyer is equal to the minimum of the value obtained by the buyer from the allocation and the buyer's budget.

Definition 5. For an allocation $x : A \times \Theta^n \to \Delta^n$, we define its liquid welfare as

$$\mathrm{LW}(x) = \sum_{i=1}^n \mathbb{E}_{\theta_i}[\min\{\mathbb{E}_{\alpha,\theta_{-i}}[w_i^T\alpha \cdot x_i(\alpha, \theta_i, \theta_{-i})], B_i\}].$$

Next, we define price of anarchy with respect to liquid welfare for pacing-based equilibria. Our definition is an instantiation of the general definition of price of anarchy introduced in Koutsoupias and Papadimitriou (2009). Before proceeding with the definition, it is worth noting an important consequence of our revenue equivalence result (Theorem 2): given an equilibrium pacing function μ, that is, a fixed point of C_0^*, the allocation under the equilibrium parameterized by μ is the same for all standard auctions. Thus, the equilibrium allocation is determined by the pacing function and is independent of the pricing rule of the standard auction, which is reflected in the following definition. For an equilibrium pacing function μ, we use x^μ to denote the allocation under the equilibrium parameterized by μ; again, this allocation is the same for all standard auctions without reserve prices.

Definition 6. The price of anarchy of pacing-based equilibria (for all standard auctions) is defined as the ratio of the worst case liquid welfare across all pacing equilibria and the optimal liquid welfare

$$\mathrm{PoA} = \frac{\inf_{\mu:\mu \in C_0^*(\mu)} \mathrm{LW}(x^\mu)}{\sup_x \mathrm{LW}(x)},$$

where the supremum in the denominator is taken over all measurable allocations x.

Because the PoA of pacing-based equilibria does not depend on the payment rule, we can work with the most convenient standard auction to prove a lower bound on the PoA, which in this case happens to be the second price auction. Azar et al. (2017) study the PoA of pure-strategy Nash equilibria of second price auctions in a non-Bayesian multi-item setting with budgets and provide a lower bound of $1/2$ for it. Unfortunately, their result hinges on the "no overbudgeting" assumption that requires the sum of equilibrium bids to be bounded above by the budget, which need not hold for pacing-based equilibria, thereby necessitating new proof ideas. Moreover, their bound may be vacuous for some parameter values because a pure-strategy Nash equilibrium is not guaranteed to exist in their setting. To get around this, they study mixed-strategy and Bayes–Nash equilibria and bound their PoA, but the lower bound they obtain for these equilibria is much worse (less than 0.02). Our model does not suffer from the problem of existence: a pure-strategy pacing-based equilibrium is always guaranteed to exist (Theorem 1). This makes the following lower bound on the PoA, which provides a worst case guarantee of $1/2$, more appealing.

Theorem 3. *The PoA of pacing-based equilibria of any standard auction is greater than or equal to 1/2.*

The proof, which is in Online Appendix D, leverages the complementary slackness condition of pacing-based equilibria to bound the PoA. Interestingly, our proof does not use a hypothetical deviation to another bidding strategy, a technique commonly found in PoA bounds (see Roughgarden et al. 2017 for a survey) and, thus, may be of independent interest.

6. Structural Properties

In this section, we show that pacing-based equilibria satisfy certain monotonicity and geometric properties related to the space of value vectors. It is worth noting that, in light of the revenue equivalence result of the preceding section, the properties established in this section hold for pacing equilibria of all standard auctions. As in Section 5, we assume that the reserve price for each item is zero, that is, $r(\alpha) = 0$ for all $\alpha \in A$. Without this assumption, similar results hold, but they become less intuitively appealing and harder to state. Moreover, we also assume that the support of G, denoted by $\delta(G)$, is a convex compact subset of $\mathbb{R}_+^{d+1}$. This assumption is made to avoid having to specify conditions on the pacing multipliers of types with probability zero of occurring. Moreover, we consider a pacing function $\mu : \Theta \to [0, \omega/B_{\min}]$ such that $\mu(w, B)$ is the unique optimal solution for the dual minimization problem for each (w, B) in the support of G, that is, $\mu(w, B) = \arg\min_{t \in [0, \omega/B_{\min}]} q^{\mu}(w, B, t)$ for all $(w, B) \in \delta(G)$. We remark that we are assuming that the best response is unique rather than the equilibrium being unique. The former can be shown to hold under fairly general conditions.

First, in Lemma 2, we showed that the pacing function associated with an SFPE is monotone in the buyer type. In particular, when the best response is unique, this result implies that $\mu(w, B)$ is nondecreasing in each component of the weight vector w and nonincreasing in the budget B. Intuitively, if the budget decreases, a buyer needs to shade bids more aggressively to meet the buyer's constraints. Alternatively, when the weight vector increases, the advertiser's paced values increase, which results in more auctions won and higher payments. Therefore, to meet the advertiser's constraints, the advertiser needs to respond by shading bids more aggressively. Furthermore, when the best response is unique, it can also be shown that μ is continuous (see Lemma 14 in the online appendix).

The next theorem further elucidates the structure imposed on μ by virtue of it corresponding to the optima of the family of dual optimization problems parameterized by (w, B). In what follows, we refer to a buyer (w, B) with $\mu(w, B) = 0$ as an *unpaced* buyer and call the buyer a *paced* buyer otherwise.

Proposition 3. *Consider a unit vector $\hat{w} \in \mathbb{R}_+^d$ and budget $B > 0$ such that $w/\|w\| = \hat{w}$ for some $(w, B) \in \delta(G)$. Then, the following statements hold:*

1. *Paced buyers with budget B and weight vectors lying along the same unit vector $\hat{w}$ have identical paced feature vectors in equilibrium. Specifically, if $(w_1, B), (w_2, B) \in \delta(G)$ with $w_1/\|w_1\| = w_2/\|w_2\| = \hat{w}$ and $\mu(w_1, B), \mu(w_2, B) > 0$, then $w_1/(1 + \mu(w_1, B)) = w_2/(1 + \mu(w_2, B))$.*
2. *Suppose there exists an unpaced buyer $(w, B) \in \delta(G)$ with $w/\|w\| = \hat{w}$ and $\mu(w, B) = 0$. Let $w_0 = \arg\max\{\|w\| \mid w \in \mathbb{R}^d;\ \mu(w, B) = 0$ and $w/\|w\| = \hat{w}\}$ be the largest unpaced weight vector along the direction $\hat{w}$. Then, all paced weight vectors get paced down to w_0, that is, $w/(1 + \mu(w, B)) = w_0$ for all $w \in \delta(G)$ with $w/\|w\| = \hat{w}$ and $\mu(w, B) > 0$.*

In combination with complementary slackness, the first part states that, in equilibrium, buyers who have the same budget, have positive pacing multipliers, and have feature vectors that are scalar multiples of each other get paced down to the same type at which they exactly spend their budget. In other words, scaling up the feature vector of a budget-constrained buyer and keeping the buyer's budget the same does not affect the equilibrium outcome. The second case of Proposition 3 addresses the directions of buyers that have a mixture of paced and unpaced buyers. In this case, there is a critical buyer type who exactly spends the buyer's budget when unpaced, and all buyer types that have weight vectors with larger norm (but the same budget) get paced down to this critical buyer type; that is, their paced weight vector equals the critical buyer type's weight vector in equilibrium. The buyer types that have a smaller norm are unpaced.

Our nonatomic model also allows us to answer the following question: keeping the competition fixed, how should an advertiser modify the advertiser's targeting criteria or ad (as captured by the weight vector) in order to maximize the advertiser's utility? This result is especially important for online display ad auctions, in which the weight vector is estimated with the goal of predicting the click-through rate and advertisers routinely modify their ads to attract more clicks. The following theorem states that the gradient w.r.t. the weight vector of the equilibrium utility of a buyer with type (w, B) is given by the expected feature vector that the buyer wins in equilibrium. This is because strong duality (Proposition 1) implies that the utility of every buyer type is given by the optimal dual value $q^\mu(w, B, \mu(w, B))$. From a practical perspective, an advertiser should focus on improving the weights of those features that have the largest average among the contexts won. It is worth noting that these quantities can be easily computed using data available to an advertiser.

Proposition 4. *Assume that A is compact. Let $\mu : \Theta \to \mathbb{R}_{\geq 0}$ be an equilibrium pacing function, that is, $\mu : \Theta \to \mathbb{R}_{\geq 0}$ such that $\mu(w, B) \in \arg\min_{t \geq 0} q^\mu(w, B, t)$ almost surely w.r.t. $(w, B) \sim G$. Then, for all $(w, B) \in \Theta$, we have $\nabla_w q^\mu(w, B, \mu(w, B)) = \mathbb{E}_{\alpha, \{\theta_i\}_{i=1}^{n-1}}[\alpha \mathbb{1}\{\beta^\mu((w, B), \alpha) \geq \beta^\mu(\theta_i, \alpha)_i \ \forall i\}]$.*

7. Analytical Example and Numerical Experiments

In this section, we illustrate our theory by providing a stylized example in which we can determine the equilibrium bidding strategies in closed form and then conduct some numerical experiments to verify our theoretical results. The purpose of the analytical example is to confirm our structural results and also help validate that our numerical procedures converge to an approximate version of the equilibrium strategies proposed in our paper.

7.1. Analytical Example

We provide an instructive (albeit stylized) example with two-dimensional feature vectors to illustrate the structural property described in Section 6. For $1 \leq a < b$, define the set of buyer types as (see the gray region in Figure 1 for a visualization of this set)

$$\Theta := \left\{ (w, B) \in \mathbb{R}^2_{\geq 0} \times \mathbb{R}_+ \,\middle|\, a \leq \|w\| \leq b,\ B = \frac{2\|w\| - w_1 - w_2}{\pi \|w\|} \right\}.$$

In this example, weight vectors lie in the intersection of a disk with the nonnegative quadrant. Observe that all buyer types whose weight vectors are colinear (i.e., they lie along the same unit vector) have identical budgets. Let the number of buyers in the auction be $n = 2$. Moreover, define the set of item types as the two standard basis vectors $A := \{e_1, e_2\}$. Finally, let G (distribution over buyer types) and F (distribution over item types) be the uniform distribution on Θ and A, respectively. Because A is discrete and F does not have a density, this example does not satisfy the assumptions we make in our model. Nonetheless, in the next claim, we show that not only does a pacing equilibrium exist, but we can also state it in closed form. The proof of the claim can be found in Online Appendix F.

Claim 1. *The pacing function $\mu : \Theta \to \mathbb{R}$ defined as $\mu(w, B) = \|w\| - 1$ for all $(w, B) \in \Theta$ is an equilibrium, that is, β^μ as given in Definition 1, is an SFPE.*

Because $H^\mu_\alpha(\cdot)$ is a strictly increasing function for all $\alpha \in A$, it is easy to check that $\mu(w, B)$ is the unique optimal to the dual optimization problem $\min_{t \in [0, \omega/B_{\min}]} q(\mu, w, B, t)$ for all $(w, B) \in \delta(G)$. Therefore, this example falls under the purview of part 1 of Proposition 3. As expected, conforming to Proposition 3, the buyers whose weight vectors are colinear get paced down to the same point on the unit arc as shown in Figure 1.

7.2. Numerical Experiments

We now describe the simulation-based experiments we conducted to verify our theoretical results. As is necessitated by computer simulations, we studied a discretized version of our problem in these experiments. More precisely, in our experiments, we used discrete approximations to the buyer type distribution G and item type

Figure 1. (Color online) Example from Section 7.1 with $a = 2, b = 3$

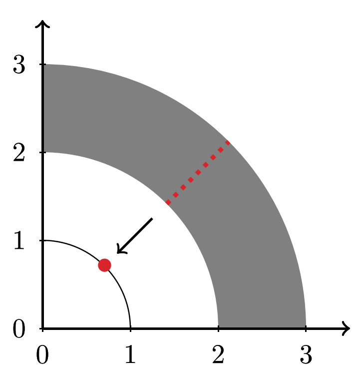

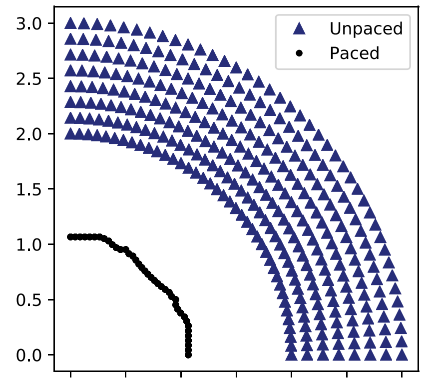

Notes. The unpaced and paced buyer weight vectors are uniformly distributed in the gray (triangle) and black (circle) regions, respectively. Each plot shows the distribution of two-dimensional buyer weight vectors. The weight vectors before pacing are depicted in gray (triangles), and the paced weight vectors are depicted in black (circles). The left plot shows the theoretical results of Section 7.1. In the left plot, the buyer weight vectors lying on the dotted line get paced down to the point. The right plot shows the results of best response iteration on the corresponding discretized problem.

distribution F. Moreover, for all item types α, we set the reserve price $r(\alpha) = 0$. One of the primary objectives of our simulations is to demonstrate that, despite the discretization, a buyer type can obtain the buyer's optimal bidding strategy by finding the optimal solution to the dual problem as our theory suggests. In other words, to compute an equilibrium, it suffices to best respond in the dual space, which has the advantage of being much simpler than the primal space. To do so, for each discretized instance, we run best response dynamics in the dual space by iterating over buyer types: computing each buyer type's optimal dual solution, keeping everyone else's pacing-based strategy fixed, and then using this optimal dual solution to determine the buyer's pacing-based bidding strategy. This approach is not guaranteed to converge. In fact, because of the discretization, strong duality may fail to hold, and a pure strategy equilibrium may not even exist. Nevertheless, despite the lack of theoretical guarantees, our experiments demonstrate that our analytical results and the dual best-response algorithm they inspire continue to work well in discrete settings.

As a first step and to validate our best response dynamics, we ran the algorithm on the discrete approximation of the example discussed in Section 7.1, for which we already analytically determined a pacing equilibrium in Claim 1. The problem was discretized by picking 320 points lying in the set of buyer types Θ defined in Section 7.1. In Figure 1, we provide plots for the case when $a = 2, b = 3$. We see that the theoretical predictions from Claim 1 are replicated almost exactly by the solution computed by the best response iteration on the discretized problem. Moreover, colinear buyer types converge to the same paced-type vector, thereby validating Proposition 3.

We conducted experiments to verify the structural properties described in Proposition 3. Here, we consider instances with $n = 3$ buyers per auction, $d = 2$ features, the buyer type distribution G given by the uniform distribution on $(1,2) \times (1,2) \times \{0.6\}$, and the item type distribution F given by the uniform distribution on the one-dimensional simplex $\{(x, y) | x, y \geq 0; x + y = 1\}$. These were discretized taking a uniform grid with 10 points along each dimension. The results are portrayed in Figure 2. The structural properties discussed in Proposition 3 are clearly evident in Figure 2. In this scenario, the buyer types are uniformly distributed on $(1,2) \times (1,2) \times \{0.6\}$, and as a consequence, all buyers have identical budgets equal to 0.6. At equilibrium, it can be seen that the colinear buyer types (i.e., buyers whose weight vectors w are colinear) who have a positive multiplier get paced down to the critical buyer type who exactly spends the buyer's budget. Moreover, at equilibrium, the boundary that separates the paced buyer types from the unpaced buyer types—the curve in which the critical buyer types lie—can be clearly observed in the left-hand plot in Figure 2. Finally, we constructed random discrete instances by uniformly sampling 50 buyer weight vectors and 20 item feature vectors from the square $(1,2) \times (1,2)$ and setting the number of buyers to be $n = 3$ and the budget of all buyer types to be $B = 2$. We found that our dual-based dynamics always converged within 250 iterations to pacing-based bidding strategies that, on average, were within 2.5% of the utility-maximizing budget feasible bidding strategy.

8. Conclusion and Future Work

This paper introduces a natural contextual valuation model and characterizes the equilibrium bidding behavior of budget-constrained buyers in first price auctions in this model. We extend this result to other standard auctions and establish revenue equivalence among them. Because of the extensive focus on second price auctions, previous work endorses bid pacing as the framework of choice for budget management in the presence of strategic buyers. Our results suggest that value pacing, which coincides with bid pacing in second price auctions, is an appropriate framework to manage budgets across all standard auctions.

An important open question we leave unanswered is that of optimizing the reserve prices to maximize seller revenue under equilibrium bidding. In general, optimizing under equilibrium constraints is challenging, so it is interesting to explore whether our model possesses additional structure that allows for tractability. Another related

Figure 2. (Color online) Numerical Experiment from Section 7.2

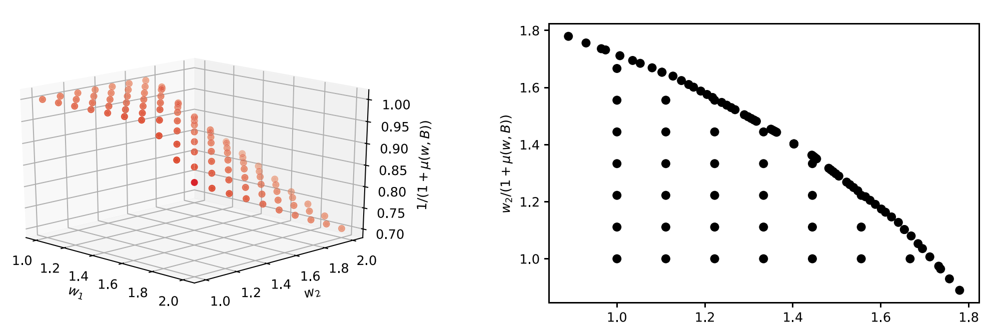

Notes. The left plot depicts how the multiplicative shading factor $1/(1 + \mu(w, B))$ varies with buyer weight vector w (budget $B = 0.6$ is the same for every buyer type). On the right, we plot the paced weight vectors of the buyer types.

question is that of characterizing the revenue-optimal mechanism for our model. Our contextual-value model can capture multi-item auctions with additive valuations as a special case (by interpreting each context as a different item), which is a notoriously hard setting for revenue maximization even in the absence of budget constraints. Investigating dynamics in first price auctions with strategic budget-constrained buyers is another interesting open direction worth exploring. We also leave open the question of efficient computation of the pacing-based equilibria discussed in this paper. Addressing this question will likely require choosing a suitable method of discretization and tie breaking without which equilibrium existence may not be guaranteed (see, e.g., Babaioff et al. 2021, Conitzer et al. 2022a). Finally, another interesting research direction is to develop conditions that guarantee uniqueness of an equilibrium. In light of recent results by Conitzer et al. (2022a), we conjecture that, without further assumptions, the equilibrium would generally not be unique.

Endnotes

[1] See https://www.emarketer.com/content/us-digital-ad-spending-2019.

[2] See https://www.blog.google/products/admanager/rolling-out-first-price-auctions-google-ad-manager-partners/.

[3] See https://www.mopub.com/en/blog/first-price-auction.

[4] The Google Ads Help page defines "average daily budget" at https://support.google.com/google-ads/answer/6312?hl=en.

[5] See, for example, https://www.blog.google/products/admanager/rolling-out-first-price-auctions-google-ad-manager-partners/.

[6] We use the term "value pacing–based" strategies to differentiate it from bid pacing/shading, which is previously studied in the context of truthful auctions (Borgs et al. 2007; Balseiro et al. 2015, 2021; Conitzer et al. 2022a, b).

References

Abhishek V, Hosanagar K (2013) Optimal bidding in multi-item multislot sponsored search auctions. *Oper. Res.* 61(4):855–873.

Aggarwal G, Badanidiyuru A, Mehta A (2019) Autobidding with constraints. Caragiannis I, Mirrokni VS, Nikolova E, eds. *Web Internet Econom. 15th Internat. Conf. Proc.* (Springer, Cham, Switzerland), 17–30.

Aliprantis CD, Border KC (2006) *Infinite Dimensional Analysis: A Hitchhiker's Guide* (Springer Science & Business Media).

Azar Y, Feldman M, Gravin N, Roytman A (2017) Liquid price of anarchy. *Internat. Sympos. Algorithmic Game Theory* (Springer, Cham, Switzerland), 3–15.

Babaioff M, Cole R, Hartline JD, Immorlica N, Lucier B (2021) Non-quasi-linear agents in quasi-linear mechanisms (extended abstract). Lee JR, ed. *12th Innovations Theoretical Comput. Sci. Conf.* (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany), 84:1–84:1.

Balseiro S, Kim A, Mahdian M, Mirrokni V (2021) Budget-management strategies in repeated auctions. *Oper. Res.* 69(3):859–876.

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

Balseiro SR, Besbes O, Weintraub GY (2015) Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Sci.* 61(4):864–884.

Bertsekas DP, Hager WW, Mangasarian OL (1998) *Nonlinear Programming* (Athena Scientific, Belmont, MA).

Bigler J (2021) Rolling out first price auctions to Google Ad Manager partners. Accessed February 17, 2021, https://www.blog.google/products/admanager/rolling-out-first-price-auctions-google-ad-manager-partners/.

Borgs C, Chayes J, Immorlica N, Jain K, Etesami O, Mahdian M (2007) Dynamics of bid optimization in online advertisement auctions. Williamson CL, Zurko ME, Patel-Schneider PF, Shenoy PJ, eds. *Proc. 16th Internat. Conf. World Wide Web* (ACM, New York), 531–540.

Che Y-K, Gale I (1998) Standard auctions with financially constrained bidders. *Rev. Econom. Stud.* 65(1):1–21.

Chen X, Kroer C, Kumar R (2021a) The complexity of pacing for second-price auctions. Biró P, Chawla S, Echenique F, eds. *22nd ACM Conf. Econom. Comput.* (ACM, New York), 318.

Chen X, Owen Z, Pixton C, Simchi-Levi D (2021b) A statistical learning approach to personalization in revenue management. *Management Sci.* 68(3):1923–1937.

Ciocan DF, Iyer K (2021) Tractable equilibria in sponsored search with endogenous budgets. *Oper. Res.* 69(1):227–244.

Conitzer V, Kroer C, Sodomka E, Stier Moses NE (2022a) Multiplicative pacing equilibria in auction markets. *Oper. Res.* 70(2): 963–989.

Conitzer V, Kroer C, Panigrahi D, Schrijvers O, Sodomka E, Stier-Moses NE, Wilkens C (2022b) Pacing equilibrium in first-price auction markets. *Management Sci.* 68(12): 8515–8535.

Dobzinski S, Leme RP (2014) Efficiency guarantees in auctions with budgets. Esparza J, Fraigniaud P, Husfeldt T, Koutsoupias E, eds. *Internat. Colloquium Automata Languages Programming* (Springer, Berlin), 392–404.

Evans LC, Gariepy RF (2015) *Measure Theory and Fine Properties of Functions* (CRC Press, New York).

Goel G, Mirrokni V, Leme RP (2015) Polyhedral clinching auctions and the AdWords polytope. *J. ACM* 62(3):1–27.

Goke S, Weintraub GY, Mastromonaco R, Seljan S (2021) Learning new auction format by bidders in internet display ad auctions. Preprint, submitted October 20, https://arxiv.org/abs/2007.00514.

Golrezaei N, Javanmard A, Mirrokni V (2021) Dynamic incentive-aware learning: Robust pricing in contextual auctions. *Oper. Res.* 69(1):297–314.

Gummadi R, Key P, Proutiere A (2011) Repeated auctions under budget constraints: Optimal bidding strategies and equilibria. *49th Annual Allerton Conference on Communication, Control, and Computing* (IEEE, Piscataway, NJ), 588.

Iyer K, Johari R, Sundararajan M (2014) Mean field equilibria of dynamic auctions with learning. *Management Sci.* 60(12):2949–2970.

Kotowski MH (2020) First-price auctions with budget constraints. *Theoretical Econom.* 15(1):199–237.

Koutsoupias E, Papadimitriou C (2009) Worst-case equilibria. *Comput. Sci. Rev.* 3(2): 65–69.

Krishna V (2009) *Auction Theory* (Academic Press, Cambridge, MA).

Kroer C, Peysakhovich A, Sodomka E, Stier-Moses NE (2022) Computing large market equilibria using abstractions. *Oper. Res.* 70(1):329–351.

Langford J, Zhang T (2007) The epoch-greedy algorithm for contextual multi-armed bandits. Platt J, Koller D, Singer Y, Roweis S, eds. *Adv. Neural Inform. Processing Systems* vol 20, (Curran Associates, Inc., Red Hook, NY).

Li L, Chu W, Langford J, Schapire RE (2010) A contextual-bandit approach to personalized news article recommendation. Rappa M, Jones P, Freire J, Chakrabarti S, eds. *Proc. 19th Internat. Conf. World Wide Web* (ACM, New York), 661–670.

Lobel I, Leme RP, Vladu A (2018) Multidimensional binary search for contextual decision-making. *Oper. Res.* 66(5):1346–1361.

McMahan HB, Holt G, Sculley D, Young M, Ebner D, Grady J, Nie L, Phillips T, Davydov E, Golovin D, et al. (2013) Ad click prediction: A view from the trenches. Dhillon IS, Koren Y, Ghani

R, Senator TE, Bradley P, Parekh R, He J, Grossman RL, Uthurusamy R, eds. *Proc. 19th ACM SIGKDD Internat. Conf. Knowledge Discovery Data Mining* (ACM, New York), 1222–1230.

Mehta A (2013) Online matching and ad allocation. *Found. Trends Theor. Comput. Sci.* 8(4):265–368.

Pai MM, Vohra R (2014) Optimal auctions with financially constrained buyers. *J. Econom. Theory* 150:383–425.

Roughgarden T, Syrgkanis V, Tardos E (2017) The price of anarchy in auctions. *J. Artificial Intelligence Res.* 59(1):59–101.

Schmeidler D (1973) Equilibrium points of nonatomic games. *J. Statist. Phys.* 7(4):295–300.

Talluri KT, Van Ryzin GJ (2006) *The Theory and Practice of Revenue Management*, vol. 68 (Springer Science & Business Media, Berlin).