



Mathematics of Operations Research

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

The Complexity of Pacing for Second-Price Auctions

Xi Chen, Christian Kroer, Rachitesh Kumar

To cite this article:

Xi Chen, Christian Kroer, Rachitesh Kumar (2024) The Complexity of Pacing for Second-Price Auctions. Mathematics of Operations Research 49(4):2109-2135. <https://doi.org/10.1287/moor.2022.0009>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2023, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

The Complexity of Pacing for Second-Price Auctions

Xi Chen,^a Christian Kroer,^a Rachitesh Kumar^{a,*}

^aColumbia University, New York, New York 10027

*Corresponding author

Contact: xichen@cs.columbia.edu (XC); ck2945@columbia.edu, <https://orcid.org/0000-0002-9009-8683> (CK); rk3068@columbia.edu, <https://orcid.org/0000-0001-7803-6861> (RK)

Received: January 6, 2022

Revised: December 25, 2022; August 9, 2023

Accepted: September 10, 2023

Published Online in Articles in Advance:
October 17, 2023

MSC2020 Subject Classifications: Primary:
91B26, 68Q25

<https://doi.org/10.1287/moor.2022.0009>

Copyright: © 2023 INFORMS

Abstract. Budget constraints are ubiquitous in online advertisement auctions. To manage these constraints and smooth out the expenditure across auctions, the bidders (or the platform on behalf of them) often employ pacing: each bidder is assigned a pacing multiplier between zero and one, and her bid on each item is multiplicatively scaled down by the pacing multiplier. This naturally gives rise to a game in which each bidder strategically selects a multiplier. The appropriate notion of equilibrium in this game is known as a pacing equilibrium. In this work, we show that the problem of finding an approximate pacing equilibrium is PPAD-complete for second-price auctions. This resolves an open question of Conitzer et al. [Conitzer V, Kroer C, Sodomka E, Stier-Moses NE (2022a) Multiplicative pacing equilibria in auction markets. *Oper. Res.* 70(2):963–989]. As a consequence of our hardness result, we show that the tâtonnement-style budget-management dynamics introduced by Borgs et al. [Borgs C, Chayes J, Immorlica N, Jain K, Etesami O, Mahdian M (2007) Dynamics of bid optimization in online advertisement auctions. *Proc. 16th Internat. Conf. World Wide Web* (ACM, New York), 531–540] are unlikely to converge efficiently for repeated second-price auctions. This disproves a conjecture by Borgs et al. [Borgs C, Chayes J, Immorlica N, Jain K, Etesami O, Mahdian M (2007) Dynamics of bid optimization in online advertisement auctions. *Proc. 16th Internat. Conf. World Wide Web* (ACM, New York), 531–540], under the assumption that the complexity class PPAD is not equal to P. Our hardness result also implies the existence of a refinement of supply-aware market equilibria which is hard to compute with simple linear utilities.

Funding: This work was supported by National Science Foundation (CCF-1703925, IIS-1838154).

Keywords: auctions • budget constraints • pacing • market equilibrium • computational complexity • PPAD

1. Introduction

Online auctions are a mainstay of the internet advertising industry. Whenever a user visits a web page or searches for a keyword, interested advertisers participate in an auction to win the opportunity to promote their content to the user. Advertisers typically participate in thousands of these online ad auctions every day and are often budget constrained, which makes budget management a crucial component of online advertising. This paper is concerned with a specific method of budget management in auctions: *pacing* (also known as multiplicative pacing), which has found use at platforms such as Facebook, where pacing is routinely employed as one of the ways to manage budgets on behalf of advertisers.¹

Pacing involves multiplicatively scaling down bids of advertisers in order to ensure a smooth depletion of their budgets over the entire advertising campaign, which is comprised of a large number of individual auctions. Consider the setting in which a group of buyers (advertisers) participate in a series of independent second-price auctions for a collection of items (the opportunity to display an ad to a user). If all buyers bid their values² in every auction, they might all deplete their budgets before the last auction. As a remedy, pacing associates a *pacing multiplier* to each buyer, which lies between zero and one, such that each buyer bids her value scaled down by her pacing multiplier. The pacing multiplier is strictly smaller than one only if the buyer would deplete her budget by bidding her true value in each auction.

Pacing has the desirable property that, if we fix the bids of competing buyers, then pacing allows a buyer to win the items which provide the best return on investment (ratio of value to price) subject to her budget constraint. In a recent work, Balseiro and Gur [4] exploit this property to prove the optimality of pacing for budget management: For a budget-constrained buyer who repeatedly participates in second-price auctions for which her values are drawn i.i.d. from some distribution, the optimal bidding strategy is to use pacing, both when the bids of the adversary are stochastic or adversarial. In other words, when considering the problem of bidding under budget constraints from the perspective of a single buyer, pacing is provably the best strategy to use.

In this paper, we study the situation where *every* buyer uses pacing to attempt to bid optimally in second-price auctions. We prove that, unfortunately, if every buyer tries to bid optimally through pacing, then the resulting dynamics are *not* likely to converge efficiently to an equilibrium. We do so by investigating the computational complexity of finding an equilibrium of the game in which each buyer's strategy involves selecting a pacing multiplier, called a *second-price pacing game*. The natural notion of equilibrium in this game is the *pacing equilibrium* (see Definition 1), which was introduced and shown to always exist by Conitzer et al. [18]. Conitzer et al. [18] also studied the computation of pacing equilibria by developing mixed-integer programming methods and applied them to real-world auction data, but found them plagued with poor scalability. They went on to conjecture that computing a pacing equilibrium could be PPAD-complete.³

Our paper resolves this open problem: it is indeed a PPAD-complete problem, and this holds for a broad class of approximate versions of the problem as well. This provides mathematical support for the repeatedly observed empirical fact that pacing-based bidding strategies often do not converge quickly, and the resulting equilibria seem hard to compute. These two facts are evidenced by the lack of efficient dynamics and efficient algorithms for computing equilibria, despite the significant attention pacing has received for more than a decade (see Section 1.4). Through our result, we show that multibuyer pacing is fundamentally intractable and this lack of efficient dynamics/algorithms is likely here to stay. Before delving deeper, we provide a short primer on PPAD for those unfamiliar with it. This can be safely skipped by any reader already familiar with the topic.

Like the well-known complexity class NP, PPAD (*Polynomial Parity Argument in a Directed graph*, introduced by Papadimitriou [38]) is a collection of computational problems. As with the definition of NP-hardness and NP-completeness, a problem is said to be PPAD-hard if it is at least as hard as every problem in PPAD; a problem is said to be PPAD-complete if it is contained in PPAD and is PPAD-hard. The analogy to NP extends further: the PPAD-hardness of a problem can be established by providing a polynomial-time reduction *from* a problem already known to be PPAD-hard. One of the quintessential PPAD-complete problems, and the one we will employ in our reductions, is that of computing a Nash equilibrium of a bimatrix game (Chen and Deng [11], Daskalakis et al. [20]). The Nash equilibrium problem has been studied extensively for decades and yet, despite much effort, no polynomial-time algorithm is known for it. Moreover, a recent spate of results showed that it is hard to solve, assuming certain strong cryptographic assumptions (Bitansky et al. [8], Choudhuri et al. [17], Garg et al. [29], Hubacek and Yogo [32], Rosen et al. [39]). This has motivated the conjecture that PPAD-hard problems cannot be solved efficiently. In this paper, we show that the problem of finding a pacing equilibrium is PPAD-hard. This shows that computing a pacing equilibrium is hard, unless all problems in PPAD can be solved efficiently.

On the other hand, showing that a problem is in PPAD amounts to giving a polynomial-time reduction *to* a problem in PPAD. For this purpose we will avail ourselves of the fact that the algorithmic version of Sperner's lemma is known to be in PPAD (Chen and Deng [12], Papadimitriou [38]), and reduce the problem of finding a pacing equilibrium to it. We refer the interested reader to Goldberg [31] and Chapter 4 of Roughgarden [40] for a survey of PPAD and its complete problems.

1.1. Main Contributions

We first prove that finding a pacing equilibrium is in PPAD. In particular, this implies that, when values and budgets of buyers are rational in the game, there always exists a pacing equilibrium in which every entry is rational and can be written using polynomially many bits. (In contrast, the existence proof of Conitzer et al. [18] uses a convergence argument, from which it is not clear whether an equilibrium with rational entries always exists.)

Theorem 1. *Finding a pacing equilibrium in a second-price pacing game is in PPAD.*

Next we show that the problem of finding an *approximate* pacing equilibrium is PPAD-hard. Our notion of approximation relaxes the definition of (exact) pacing equilibria in two ways: (i) buyers who bid *close* to (but not necessarily exactly equal to) the highest bid may also win fractions of an item; (ii) each buyer either spends *most* of her budget, or her pacing multiplier is *close* to one. We use two parameters δ and γ to capture these two relaxations quantitatively and such a solution is called a (δ, γ) -*approximate pacing equilibrium* (see Definition 2).

Theorem 2. *For any constant $c > 0$, finding a (δ, γ) -approximate pacing equilibrium in a second-price pacing game with n players is PPAD-hard when $\delta = \gamma = 1/n^c$.*

Note that, by virtue of being a relaxation, finding an approximate pacing equilibrium is in PPAD as a direct consequence of Theorem 1. Similarly, the PPAD-hardness of finding an exact pacing equilibrium follows from Theorem 2. Therefore, both problems of finding an exact and an approximate pacing equilibrium are complete in PPAD. To the best of our knowledge, our results are the first PPAD-completeness results for budget-management in second-price auction systems such as those applied in large-scale internet advertising.

1.1.1. Implications. Our hardness result has implications for Borgs et al. [9], in which the authors studied dynamic first-price and second-price auctions with budgets. They proved that pacing combined with perturbations can lead to efficient convergence of bidding dynamics under first-price auctions. They conjectured a similar convergence in the analogous second-price setting and provided experimental support for it. Our definition of approximate pacing equilibria (Definition 2) is able to capture their random-perturbation model, thereby bringing it under the purview of our hardness result Theorem 2: if such a convergence occurs in the second-price case, then it must do so inefficiently assuming PPAD does not have polynomial-time algorithms (see Section 3.2). Moreover, since our model also admits a stochastic interpretation, if all of the buyers employ some pacing algorithm for repeated second-price auctions with correlated value distributions and global budget constraints, then the resulting dynamics will not always converge efficiently to an equilibrium, assuming PPAD does not have polynomial-time algorithms. In particular, this statement applies to the pacing algorithm given by Balseiro and Gur [4], which is an optimal bidding algorithm for a single budget-constrained buyer under both adversarial and independent-stochastic competition. Informally, the central message here is that, when multiple budget-constrained buyers bid in a way that is optimal for them individually, the resulting dynamics will not in general stabilize to an equilibrium.

Furthermore, due to connections between pacing equilibria and *supply-aware* market equilibria (Conitzer et al. [18]) with linear utilities, our PPAD-hardness result has novel consequences when interpreted in the language of market equilibria: our result shows that a natural refinement of supply-aware market equilibria with linear utilities is PPAD-hard (finding one with prices corresponding to second-price auctions).

1.2. Techniques Used

We prove the PPAD-hardness of finding approximate pacing equilibria (Theorem 2) by giving a reduction from the problem of finding an ϵ -well-supported Nash equilibrium in win-lose bimatrix games. The second-price rule plays an important role in this reduction. Consider an item with two interested buyers, one of which has a much higher value than the other, so much so that she always wins the good in any pacing equilibrium. Then, the payment made by this buyer on this item is determined by the bid of the lower-valued buyer, which is equal to her value times her multiplier. This allows us to construct gadgets which capture Nash equilibria of any bimatrix game with pacing multipliers, by using the second-price rule to account for the expected cost of each action of a player with respect to the other player's mixed strategy. A complicating factor in our proof is that pacing multipliers are always positive, whereas some actions are played with probability zero in a Nash equilibrium. To address this issue, we construct our gadgets such that they have a discontinuous behavior: there is a baseline pacing amount which corresponds to playing the corresponding action with probability zero, and only larger pacing values correspond to probabilities.

To prove the PPAD-membership of finding a pacing equilibrium (Theorem 1), we reduce the problem to the algorithmic version of Sperner's Lemma. A direct reduction proves challenging due to the discontinuous way in which the allocation of an item varies with pacing multipliers: In a pacing equilibrium, an item can only be assigned to buyers whose bids are *exactly* equal to the highest bid. Similar issues were encountered in PPAD-membership proofs for market equilibrium computation (Vazirani and Yannakakis [41]). For this reason, we start by proving the PPAD-membership of finding approximate pacing equilibria, in which items can be allocated smoothly. Then we bootstrap this result to show the PPAD-membership of exact pacing equilibrium in two steps. The first step starts with an approximate pacing equilibrium and rounds it to obtain a pacing equilibrium in which only buyers tied for the highest bid on a good share it. We still allow the relaxation that each buyer can either spend *most* of her budget or set her pacing multiplier *close* to one. Finally, we do away with this remaining relaxation by using a LP-based technique similar to the one used in Etessami and Yannakakis [24], Vazirani and Yannakakis [41], and Filos-Ratsikas et al. [27], thereby showing the PPAD-membership of finding an exact pacing equilibrium.

1.3. Pacing in Internet Advertising

To motivate pacing equilibrium as a solution concept, this section describes how the solution concept arises in practice as part of internet advertising platforms such as those operated by, for example, Facebook, Google, or Twitter. As discussed previously, pacing equilibrium may arise through individual buyers optimizing their spending due to their budget constraint. A second reason that pacing equilibrium is of practical interest is due to *proxy bidders*. When an advertiser starts a campaign, they often specify only a small set of parameters: their value for a click (or some other notion of converting an ad into value, say a video view), their budget, and their *targeting criteria* which specify the subset of users they are interested in (e.g., “people who surf” if the ad is for surfboards). Then, whenever an auction is run to determine which ads to show to a given user, the bid from a given advertiser is submitted by the proxy bidder acting on behalf of that advertiser. The proxy bidder calculates the value that advertiser i has for being shown to the user in auction j as $v_{ij} = v_i \cdot CTR_{ij}$, where v_i is the value per click and CTR_{ij} is the estimated probability that the user will click on the ad. If there were no budgets, then the proxy bidder should submit the bid v_{ij} , due to the

truthfulness of the second-price auction. But in the presence of budgets, this may negatively affect the overall utility achieved by the advertiser, since they will run out of budget well before the campaign ends, and thus miss out on later strong bang-per-buck opportunities.

To address their budget constraints, the advertisers are typically offered one or more options for budget-management strategies that can be employed by the proxy bidders. Pacing as defined in this paper, via multiplicative bid scaling, is offered by Facebook by default (Conitzer et al. [18], Facebook [25]), and it is also offered on other platforms. Intuitively speaking, the proxy bidder attempts to choose a pacing multiplier which will spend the advertiser's budget evenly across the campaign length. To ensure that this will happen, the pacing multiplier is adapted over time using a control algorithm: the algorithm will adjust the pacing multiplier up or down depending on whether the proxy bidder is currently under or overspending. Since we do not consider the online aspect of the problem, the pacing equilibrium solution concept that we study corresponds to the steady-state that this adaptive process would ideally arrive at (this is analogous to what was done by Balseiro et al. [3], Balseiro et al. [5], Conitzer et al. [18]). See Conitzer et al. [18] for a longer discussion of the pacing equilibrium model and how it relates to real-world systems.

1.4. Additional Related Work

There is a large literature on budgets in auctions, largely inspired by the internet advertising industry. Here we survey the ones most related to our paper. We start by surveying the literature on multi-item first or second-price auctions with budgets and the associated equilibrium issues there, since that is the setting we study. We briefly mention some pointers to alternative approaches and models such as mechanism design or online matching.

Balseiro et al. [5] studied budget management in second-price auctions using a *fluid mean-field* model, and showed that in this model existence is guaranteed, and closed-form solutions for equilibria are derived for certain settings (Balseiro et al. [3]). studied several different pacing mechanisms for second-price auctions, including multiplicative pacing, and showed existence results for their setting, as well as other analytical and numerical properties (Conitzer et al. [19]). studied the model of Conitzer et al. [18], but with each auction using a first-price rule. There, pacing equilibrium no longer constitutes best responses, but instead has a market equilibrium interpretation. In the first-price setting, pacing equilibria turn out to be easy to compute, due to a direct relationship to market equilibria (Babaioff et al. [2]). studied non-quasi-linear agents participating in mechanisms designed for quasi-linear agents. They studied a generalization of budget constraints where agents have a concave disutility in payment, and showed that a Nash equilibria exists which employs multiplicative scaling (Balseiro and Gur [4]). developed online learning methods for individual agents adapting their pacing multipliers over time, and showed that this converges to an equilibrium under certain stochastic independence assumptions. Assuming $PPAD \neq P$, our results can be interpreted to mean that, in the general setting which allows for correlation and discrete valuations, no dynamics can converge efficiently in the worst case (see Proposition 10 of Conitzer et al. [18] for a formal statement connecting the stochastic and deterministic settings).

An alternative approach for handling budget constraints in multi-item settings is to design a mechanism that accounts for this explicitly (see e.g., Ashlagi et al. [1], Dobzinski and Leme [22], Dobzinski et al. [23], Goel et al. [30]). Another approach to budget-constrained allocation in online advertising is to treat the problem as an online matching problem. This research was initiated by Mehta et al. [36], see e.g., Mehta [35] for a survey.

Our results are strongly related to the problem of computing market equilibria under a supply-aware model (see Section 3.2 for a discussion). There have been several PPAD-completeness results for various Fisher market models (without supply-awareness). However, these results are all for models with more complex utility functions, which give rise to the hardness (Chen and Teng [13]), and Vazirani and Yannakakis [41] showed that for additively-separable piecewise-linear concave utilities, finding an equilibrium in a Fisher market is PPAD-complete. Bei et al. [7] showed PPAD-hardness of finding market equilibria with budget-capped utilities (this is proved using a variation on the piecewise-linear utilities proof of Chen and Teng [13]). In the case of indivisible goods, Othman et al. [37] showed that finding an approximate market equilibrium is hard, even one which is guaranteed to exist (Budish [10]). For the Arrow-Debreu exchange economy, Chen et al. [15] showed that finding an equilibrium is PPAD-hard.

Finally, in addition to Nash equilibrium and market equilibrium, many interesting problems have been proven to be PPAD complete in domains like auctions (Chen et al. [14], Filos-Ratsikas et al. [28]), fair division (Deng et al. [21], Filos-Ratsikas et al. [27]), and optimization (Fearnley et al. [26]). Recently, Li and Tang [34] studied pacing-based equilibria in autobidding where each advertiser aims to maximize total value (instead of utility) subject to a Return-on-Investment (ROI) constraint. They show that finding such equilibria is PPAD hard for second-price auctions.

2. Model

We start with the definition of Second-price Pacing Games. In a *Second-price Pacing Game* (SPP game as a shorthand) $G = (n, m, (v_{ij}), (B_i))$, there are n buyers and m (indivisible) goods. Each good is sold through independent (single slot)

second-price auctions. We use $v_{ij} \geq 0$, $i \in [n]$ and $j \in [m]$, to denote the value of good j to buyer i , and $B_i > 0$ to denote the budget of buyer i . We will require (1) for each $j \in [m]$, $v_{ij} > 0$ for some $i \in [n]$, and (2) for each $i \in [n]$, $v_{ij} > 0$ for some $j \in [m]$. Each buyer i plays the game by picking a *pacing multiplier* $\alpha_i \in [0, 1]$ and then bidding $\alpha_i v_{ij}$ on good j for each $j \in [m]$.

To finish describing the game, one approach is to specify a tie-breaking rule: a rule that determines the probabilities with which a good is allocated among the highest bidders. However Conitzer et al. [18], showed that the choice of tie-breaking rule affects equilibrium existence. This motivated them to introduce an equilibrium notion called the *pacing equilibrium*, which is not concerned with any specific tie-breaking rule, but instead includes the probability distribution used to allocate each good as part of the equilibrium (see Definition 1). We will take a similar approach and work with *pacing equilibrium*, focusing on its computational aspects. It is worth pointing out that this only makes our hardness results stronger because they apply to any tie-breaking rule (such as the one used by Borgs et al. [9], which works via random perturbations; see Section 3.3 for a detailed discussion of the implications of our hardness results).

With slight abuse of notation, we will write $x_{ij} \geq 0$ to denote the *fraction* of good j allocated to buyer i , which, in our indivisible goods regime, should be interpreted to mean the probability of allocating good j to buyer i . Therefore, the allocation should always satisfy $\sum_{i \in [n]} x_{ij} \leq 1$ for all $j \in [m]$. In addition, only buyers i with the *highest* bid for good j can have $x_{ij} > 0$ and they pay for good j under the *second-price* rule.

Formally, when the buyers use pacing multipliers $\alpha = (\alpha_1, \dots, \alpha_n)$, we let $h_j(\alpha) = \max_{i \in [n]} \alpha_i v_{ij}$ denote the *highest* bid on good j and $p_j(\alpha)$ denote the *second highest bid* on good j , that is, $p_j(\alpha)$ is the second largest element among $\alpha_1 v_{1j}, \dots, \alpha_n v_{nj}$ (in particular, $p_j(\alpha) = h_j(\alpha)$ when there is a tie for the highest bid). Only buyers who bid $h_j(\alpha)$ can purchase (fractions of) good j under the price $p_j(\alpha)$. Thus, under an allocation $x = (x_{ij})$, the total payment of buyer i is given by $\sum_{j \in [m]} x_{ij} p_j(\alpha)$, which should not exceed the budget B_i of buyer i .

Next, we define the notion of *pacing equilibria* (Conitzer et al. [18]) of SPP games. A *pacing equilibrium* consists of a tuple of pacing multipliers $\alpha = (\alpha_i)$ and an allocation $x = (x_{ij})$ of goods that satisfy the two conditions described above (i.e., only buyers with the highest bid can be allocated a good and their budgets are satisfied, as captured in (a) and (c) below). In addition, we require (b) the full allocation of any good with a positive bid and (d) that there is no unnecessary pacing: if a buyer i does not spend her whole budget, then her pacing multiplier should be one. Intuitively, this makes sense because if her budget is not binding, then she should participate as if each auction is a regular second-price auction.

Definition 1 (Pacing Equilibria). Given an SPP game $G = (n, m, (v_{ij}), (B_i))$, we say (α, x) with $\alpha = (\alpha_i) \in [0, 1]^n$, $x = (x_{ij}) \in [0, 1]^{nm}$ and $\sum_{i \in [n]} x_{ij} \leq 1$ for all $j \in [m]$ is a *pacing equilibrium* if

- (a) Only buyers with the highest bid win the good: $x_{ij} > 0$ implies $\alpha_i v_{ij} = h_j(\alpha)$.
- (b) Full allocation of each good with a positive bid: $h_j(\alpha) > 0$ implies $\sum_{i \in [n]} x_{ij} = 1$.
- (c) Budgets are satisfied: $\sum_{j \in [m]} x_{ij} p_j(\alpha) \leq B_i$.
- (d) No unnecessary pacing: $\sum_{j \in [m]} x_{ij} p_j(\alpha) < B_i$ implies $\alpha_i = 1$.

We will work with an *approximate* version of *pacing equilibria* in both of our PPAD-hardness and PPAD-membership results.

In an approximate *pacing equilibrium*, we make two relaxations on (b) and (d); the two parameters used to capture these two relaxations are δ and γ , respectively.

Definition 2 (Approximate Pacing Equilibria). Given an SPP game $G = (n, m, (v_{ij}), (B_i))$ and parameters $\delta, \gamma \in [0, 1]$, we say (α, x) , with $\alpha = (\alpha_i) \in [0, 1]^n$, $x = (x_{ij}) \in [0, 1]^{nm}$ and $\sum_{i \in [n]} x_{ij} \leq 1$ for all $j \in [m]$, is a (δ, γ) -approximate *pacing equilibrium* of G if

- (a) Only buyers *close* to the highest bid win the good: $x_{ij} > 0$ implies $\alpha_i v_{ij} \geq (1 - \delta) h_j(\alpha)$.
- (b) Full allocation of each good with a positive bid: $h_j(\alpha) > 0$ implies $\sum_{i \in [n]} x_{ij} = 1$.
- (c) Budgets are satisfied: $\sum_{j \in [m]} x_{ij} p_j(\alpha) \leq B_i$.
- (d) *Not too much* unnecessary pacing: $\sum_{j \in [m]} x_{ij} p_j(\alpha) < (1 - \gamma) B_i$ implies $\alpha_i \geq 1 - \gamma$.

For convenience we will write (δ, γ) -approximate PE to denote (δ, γ) -approximate *pacing equilibrium*, and write γ -approximate PE to denote $(0, \gamma)$ -approximate PE. It is clear from the definition that when $\delta = \gamma = 0$, (δ, γ) -approximate PE captures the exact *pacing equilibria* of a SPP game.

Remark 1. We can incorporate reserve prices in our model. Definition 1 can be extended in a natural way to model the presence of reserve prices (see Definition C.1). All our results continue to hold with this extension. We refer the reader to Appendix C for a full discussion.

2.1. Connections to Dynamics, Best Response, and Nash Equilibrium

Before moving on to our results, we motivate the definition of pacing equilibrium by connecting it more concretely to practice and previous work. Consider a collection of n buyers that participate repeatedly in T second-price auctions. For each auction $t \in [T]$, the good to be sold is drawn from a collection of m possible goods, with good j being selected with probability $d_j > 0$. Moreover, suppose the value v'_{ij} that buyer i has for good j is given by $\epsilon_{ij}v_{ij}/d_j$ for some $v_{ij} \geq 0$, where ϵ_{ij} is drawn independently for each buyer-good pair from some continuous distribution supported over $[1 - \delta, 1]$. The ϵ_{ij} component of the value can also be thought of as a perturbation that arises from errors in estimating the click-through-rate (probability of a click) which is a crucial factor in determining the value of an advertiser in internet advertising. Finally, let B'_i denote the budget of buyer i , which is the maximum amount she is willing to spend over all T auctions.

Balseiro and Gur [4] prove that, if we fix the bidding strategy of the other buyers, then it is optimal for a buyer to use pacing-based strategy to bid. The optimal pacing-based algorithm of Balseiro and Gur [4] iteratively updates the pacing multiplier and satisfies the following properties: (i) If the buyer spends less than her per-period budget $B_i = B'_i/T$ in an iteration, her pacing multiplier is increased, and if the payment is greater than her per-period budget, then the multiplier is decreased; (ii) The pacing multiplier is constrained to belong to $[0, 1]$ because bidding more than the value leads to negative utility. These properties are also satisfied by the algorithm proposed by Borgs et al. [9] and forms the basis of pacing algorithms used in practice which aim to smooth the expenditure of a buyer by evenly spending the budget over all auctions, that is, aim to spend the per-period budget in each period if possible. If all of the buyers use an algorithm that satisfies these properties, the system can only stabilize when all of the buyers satisfy the no-unnecessary-pacing condition.

The no-unnecessary-pacing condition and the optimality of pacing stem from strong duality, as argued in Balseiro et al. [5] and Balseiro and Gur [4]. We provide a brief overview of their argument here. When T is large and $B'_i = \Theta(T)$, as is the case in online advertising, concentration arguments kick in and the problem of repeatedly bidding in T auctions can be interpreted as repeatedly bidding in the following single-shot game: Each buyer wishes to maximize her expected utility (value – payment) while keeping her expenditure below $B_i = B'_i/T$ in expectation over the randomness in the values (see Balseiro and Gur [4] and Balseiro et al. [5] for more details). This single-shot game captures the crux of the problem and its variants have been extensively studied in the literature (Babaioff et al. [2], Balseiro et al. [3], Balseiro and Gur [4], Balseiro et al. [6]). In fact, Balseiro and Gur [4] show that, under some fairly stringent assumptions, their algorithm efficiently converges to an approximate pacing equilibrium of this single-shot game when all of the buyers employ it. But, these assumptions require independence of values across buyers and strong monotonicity of payments as a function of the pacing multipliers, both of which are unlikely to hold in practice. As we show in this paper, if $\text{PPAD} \neq \text{P}$, then the convergence can no longer be efficient in the absence of these assumptions. In the rest of this subsection, we will restrict our focus to this single-shot game and connect it to SPP games and pacing equilibria.

Fix buyer i and let f_j denote the highest bid from buyers other than i on good j . Then, the optimization problem faced by buyer i in the single-shot game is given by

$$\begin{aligned} \max_b \quad & \sum_{j=1}^m d_j \cdot \mathbb{E}_{v'_{ij}, f_j} [(v'_{ij} - f_j) \mathbf{1}(b(j, v'_{ij}) \geq f_j)] \\ \text{s.t.} \quad & \sum_{j=1}^m d_j \cdot \mathbb{E}_{v'_{ij}, f_j} [f_j \cdot \mathbf{1}(b(j, v'_{ij}) \geq f_j)] \leq B_i \end{aligned}$$

where $b(j, \cdot)$ denotes the bidding strategy of buyer i for good j . Assume that the distribution of f_j conditioned on v'_{ij} (value of buyer i for good j) is continuous. Then, using the strong-duality argument of Balseiro et al. [5] or Balseiro et al. [6], it can be shown that strong duality holds, where the dual problem is given by

$$\begin{aligned} \min_{\mu_i \geq 0} \quad & \mu_i \cdot B + \max_b \sum_{j=1}^m d_j \cdot \mathbb{E}_{v'_{ij}, f_j} [(v'_{ij} - (1 + \mu_i)f_j) \mathbf{1}(b(j, v'_{ij}) \geq f_j)] \\ = \min_{\mu_i \geq 0} \quad & \mu_i \cdot B + (1 + \mu_i) \max_b \sum_{j=1}^m d_j \cdot \mathbb{E}_{v'_{ij}, f_j} \left[\left(\frac{v'_{ij}}{1 + \mu_i} - f_j \right) \mathbf{1}(b(j, v'_{ij}) \geq f_j) \right] \end{aligned}$$

Therefore, if $\mu_i^* \geq 0$ is the optimal dual solution, then an optimal bidding strategy for buyer i is $b(j, v'_{ij}) = v'_{ij}/(1 + \mu_i^*)$ (i.e., to pace her value with the multiplier $\alpha_i = 1/(1 + \mu_i^*)$) since it is optimal for the inner Lagrangian optimization problem over b . Note that this argument does not require other buyers to use a pacing-based strategy. Thus, it establishes that a pacing-based best response always exists.

Strong duality also implies that any optimal primal-dual solution pair satisfies complementary slackness: $\mu_i^* = 0$ if

$$\sum_{j=1}^m d_j \cdot \mathbb{E}_{v'_{ij}, f_j} [f_j \cdot \mathbf{1}(v'_{ij}/(1 + \mu_i^*) \geq f_j)] < B_i.$$

The fixed-point argument of Balseiro et al. [5] further shows that a pacing-based Nash equilibrium exists for the single-shot game where all of the buyers use pacing with multipliers $\alpha_i = 1/(1 + \mu_i)$. Moreover, if a collection of feasible dual multipliers satisfy complementary slackness and the corresponding pacing-based strategies satisfy the budget constraints, then they form a Nash equilibrium of the single-shot game described above. Now, let $\alpha_i = 1/(1 + \mu_i)$ be a collection of equilibrium pacing multipliers. Then, the complementary slackness condition for buyer i can equivalently be written as a no-unnecessary-pacing condition: $\alpha_i = 0$ if

$$\sum_{j=1}^m d_j \cdot \mathbb{E}_{v'_{ij}, f_j} [f_j \cdot \mathbf{1}(\alpha_i v'_{ij} \geq f_j)] < B_i$$

As a consequence, every pacing equilibrium of this single-shot game is also a Nash equilibrium, where we define a pacing equilibrium to be any collection of pacing multipliers that satisfy the no-unnecessary-pacing condition and satisfy the budget constraint. Even if one has no interest in duality, the no-unnecessary-pacing condition is also extremely desirable in practice when the platform manages the budget of the buyer on her behalf—it ensures that the platform bids the value of the buyer on each good unless doing so would violate her budget. Thus, as outlined above, pacing equilibrium is an important refinement of Nash equilibrium for the single-shot game in both theory and practice.

Next, we connect pacing equilibria in single-shot games to approximate pacing equilibria in SPP games. Observe that, when all of the buyers use pacing to bid, $f_j = \max_{k \neq i} \alpha_k \epsilon_{kj} v_{kj} / d_j$. Hence, the expected payment of buyer i in this single-shot game can be rewritten as

$$\mathbb{E}_{\{\epsilon_{ij}\}_{i,j}} \left[\sum_{j=1}^m \left\{ \max_{k \neq i} \alpha_k \epsilon_{kj} v_{kj} \right\} \mathbf{1} \left(\epsilon_{ij} \alpha_i v_{ij} \geq \max_{k \neq i} \epsilon_{kj} \alpha_k v_{kj} \right) \right]$$

If we ignore the perturbations ϵ_{ij} , this is exactly the payment of buyer i in the SPP game with values v_{ij} and pacing multipliers α_i . To account for the perturbations and connect the single-shot game to the SPP game, we can define a perturbed SPP game (like Borgs et al. [9]) as one in which (i) the value of buyer i for good j is given by $\epsilon_{ij} v_{ij}$; (ii) each item is sold through second-price auction; (iii) the strategy of each buyer is her pacing multiplier $\alpha_i \in [0, 1]$; (iv) ϵ_{ij} are drawn i.i.d. from some distribution with a positive density over $[1 - \delta, 1]$; (v) each buyer wishes to maximize her expected utility while satisfying her budget constraint in expectation over the perturbations ($-\infty$ utility if the budget constraint is violated). We define an approximate pacing equilibrium of this perturbed SPP game as simply a collection of budget-feasible pacing multipliers that satisfy the not-too-much-unnecessary-condition (see Appendix D). Recall that approximate pacing equilibrium of SPP games allows for arbitrary allocation between all buyers close to the highest bid, and therefore includes the allocation induced by perturbations as a special case. In Appendix D, we use this fact to show that computing a pacing equilibrium of perturbed SPP games is harder than computing an approximate pacing equilibrium in (unperturbed) SPP games, and therefore PPAD-hard due to Theorem 3.

Finally, as we make δ smaller, this perturbed SPP game gets closer to a true SPP game. Unfortunately, the duality-based existence argument of Balseiro et al. [5] and Balseiro et al. [6] breaks down when $\delta = 0$ because ties are no longer a zero-probability event. The following example shows that a pacing equilibrium may not exist in this case under the uniform tie-breaking rule.

Example 1. Consider a setting with two buyers and one good. $v_{11} = 1$, $v_{21} = v \gg 1$ and $B_1 = \infty$, $B_2 = 1/4$. Then, in any pacing equilibrium we have $\alpha_1 = 1$ because of the no-unnecessary-pacing condition. Now, if $\alpha_2 \geq 1/v$, then buyer 2 spends at least $1/2$ due to the uniform tie-breaking rule, which violates her budget. Hence, $\alpha_2 < 1/v$ and buyer two wins nothing and spends 0, thereby violating the no-unnecessary pacing condition.

Conitzer et al. [18] show that a pacing equilibrium does exist if the ties are broken carefully, which was their motivation behind making the tie-breaking rule a part of the equilibrium concept. This equilibrium tie-breaking rule can be thought of as the limiting expected allocation in the perturbed equilibrium as δ approaches zero. They also show that, in an unperturbed SPP game, if we fix the bids of other buyers and allow a buyer to pick her bids along with the fraction of each good she wants, it is a best-response for her to use pacing to bid because it allows her to win goods that yield the highest value per unit cost—using the multiplier α_i ensures that a buyer wins a good if and only if α_i times her value is greater than the second-highest bid, that is, if the value per unit cost is above $1/\alpha_i$. In the one-shot game studied by Conitzer et al. [18], pacing may not be a best response if the

m good types arrive one-by-one and a buyer can change her bid to cause other buyers to drop out of later auctions due to budget exhaustion. In the repeated auction setting that motivates our single-stage pacing game, the budget constraint is over all T auctions. Therefore, deviating in a single-stage game would not cause other advertisers to drop out of that game.⁴ Conitzer et al. [18] also provide a discussion on the undesirable properties of Nash equilibria in SPP games en route to motivating pacing equilibria as a more desirable solution concept. Nevertheless, we would like to note that our hardness result can be extended to Nash equilibria: In Appendix D, we prove that computing a Nash equilibrium of the perturbed SPP game is also PPAD-hard. We do so by showing that a minor modification of the game constructed in our hardness reduction for Theorem 3 only admits Nash equilibria that are also pacing equilibria.

3. Hardness Results

In this section we investigate the hardness of computing approximate pacing equilibria and show that the problem is PPAD-hard for second-price pacing games. Our most general result (Theorem 2) shows that the problem of finding a (δ, γ) -approximate PE in a SPP game is PPAD-hard, even when δ and γ are polynomially small in the number of players.

Our result is shown by reducing the problem of computing a Nash equilibrium in a $\{0, 1\}$ -cost bimatrix game to that of finding a (δ, γ) -approximate PE in a corresponding SPP game. Because we wish to show the result for (δ, γ) -approximate PE, we must start our reduction from such approximate PE. In order to manage the resulting approximation factors, we are forced to introduce a number of additional bookkeeping gadgets, and correspondingly work with the problem of computing ϵ -well-supported Nash equilibria of $\{0, 1\}$ -cost bimatrix games, as opposed to standard Nash equilibria. Taken together, all these facts lead to a longer proof that may obfuscate the main ideas underlying our reduction. To better highlight the key ideas in our reduction and motivate our techniques, we are going to start by proving that finding an *exact* pacing equilibrium in a SPP game is PPAD-hard, by showing a reduction from the problem of finding an exact Nash equilibrium in a $\{0, 1\}$ -cost bimatrix game.

3.1. Hardness of Finding Exact Pacing Equilibria

Our reduction will be from the problem of computing a Nash equilibrium in a $\{0, 1\}$ -cost bimatrix game. Let Δ_n denote the set of probability distributions over $[n]$. The input of the bimatrix problem is a pair of cost matrices $A, B \in \{0, 1\}^{n \times n}$ and the goal is to find a Nash equilibrium $(x, y) \in \Delta_n \times \Delta_n$, meaning that x minimizes cost given y , that is, $x^T A y \leq \hat{x}^T A y$ for all $\hat{x} \in \Delta_n$, and similarly y minimizes cost given x , that is, $x^T B y \leq x^T B \hat{y}$ for all $\hat{y} \in \Delta_n$. Equivalently, (x, y) is a Nash equilibrium if $x_i > 0$ for any $i \in [n]$ implies that $\sum_j A_{ij} y_j \leq \sum_j A_{kj} y_j$ for all $k \in [n]$, and $y_j > 0$ for any $j \in [n]$ implies that $\sum_i x_i B_{ij} \leq \sum_i x_i B_{ik}$ for all $k \in [n]$. This problem is known to be PPAD-complete (Chen et al. [16]).

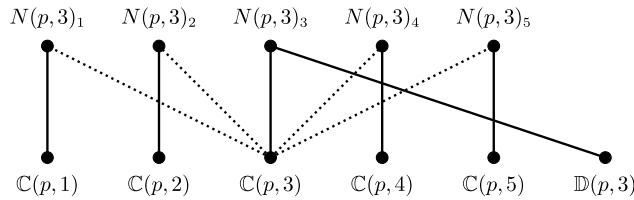
Given a $\{0, 1\}$ -cost bimatrix game (A, B) with $A, B \in \{0, 1\}^{n \times n}$, we would like to construct an SPP game G in time polynomial in n , such that every exact PE of G can be mapped back to a Nash equilibrium of the bimatrix game (A, B) in polynomial time.

Before proceeding further, we informally describe some important aspects of the construction to provide some intuition. First, in the SPP game G , we will encode the pair (x, y) of mixed strategies in Δ_n using pacing multipliers. For each player $p \in \{1, 2\}$ in the bimatrix game (A, B) and each (pure) strategy $s \in [n]$, there will be a corresponding buyer $\mathbb{C}(p, s)$ in the SPP game G , whose pacing multiplier $\alpha(\mathbb{C}(p, s))$ will be used to encode the probability with which player p plays strategy s in the bimatrix game (A, B) . For now, take x to be the distribution obtained by normalizing $\alpha(\mathbb{C}(1, s))$, that is, $x_t = \alpha(\mathbb{C}(1, t)) / \sum_s \alpha(\mathbb{C}(1, s))$, and define y similarly using $\alpha(\mathbb{C}(2, s))$; we will discuss the issues with this proposal and ways to fix them momentarily.

Second, in order to capture the best response condition of Nash equilibria, we need to encode the cost borne by player $p \in \{1, 2\}$ when playing a given strategy $s \in [n]$ against the mixed strategy of the other player. For simplicity, let us focus on $p = 1$. We will create a set of n expenditure goods $E(1, s)_1, \dots, E(1, s)_n$ for each pure strategy s of player 1. We will set buyer $\mathbb{C}(1, s)$'s value at 1 for each of the expenditure goods $E(1, s)_1, \dots, E(1, s)_n$. Additionally, each buyer $\mathbb{C}(2, t)$ will value $E(1, s)_t$ at $v A_{st}$, where $v = 1/(16n)$ is set to be so small that $\mathbb{C}(1, s)$ always wins all the goods $E(1, s)_1, \dots, E(1, s)_n$ under any PE of G . This means that, in any PE with multipliers $\alpha(\mathbb{C}(p, s))$, buyer $\mathbb{C}(1, s)$ pays a total of $v \sum_t \alpha(\mathbb{C}(2, t)) A_{st}$ for the expenditure goods $E(1, s)_1, \dots, E(1, s)_n$, which captures player 1's cost for playing strategy s in (A, B) , when player 2 uses the mixed strategy that plays each t with probability defined by $\alpha(\mathbb{C}(2, t))$ after normalization.

Finally we need to make sure that the best response condition of Nash equilibria holds for a strategy pair (x, y) obtained from multipliers $\alpha(\mathbb{C}(p, s))$ in any PE of G , that is, only best-response strategies are played with positive probability. This poses a challenge because pacing multipliers are never zero in a pacing equilibrium, so we can't use them directly to encode probabilities in x and y (which need to be zero for strategies which are not best responses). To

Figure 1. Normalization goods for $p \in \{1, 2\}$ and $s = 3$, when $n = 5$.



Notes. A buyer having a nonzero value for a good is represented by a line connecting the two. Solid line denotes a value of 1 and dotted line denotes a value of 2.

get around this issue, we will use thresholds to encode entries of (x, y) using $\alpha(\mathbb{C}(p, s))$. More formally, we add a *threshold buyer* and a set of *threshold goods* to G to make sure that $\alpha(\mathbb{C}(p, s)) \geq 1/2$ in any PE of G . This allows us to encode x by normalizing $\alpha(\mathbb{C}(1, s)) - 1/2$ and y by normalizing $\alpha(\mathbb{C}(2, s)) - 1/2$. The most challenging part of the construction is to have buyers/goods work together to ensure that both $\alpha(\mathbb{C}(1, s)) - 1/2$ and $\alpha(\mathbb{C}(2, s)) - 1/2$, $s \in [n]$, are not identically zero. We accomplish this by creating a set of *normalization goods* for each buyer $\mathbb{C}(p, s)$, with the property that each buyer $\mathbb{C}(p, s)$ spends approximately $\sum_{t=1}^n \alpha(\mathbb{C}(p, t))$ on her normalization goods. This, in combination with a carefully chosen budget and the ‘No unnecessary pacing’ condition, ensures that $\{\alpha(\mathbb{C}(p, s)) - 1/2\}_s$ are not identically zero. Then, we can follow the plan described in the last paragraph to encode the cost of player p playing s using the expenditure of buyer $\mathbb{C}(p, s)$ on $E(p, s)_1, \dots, E(p, s)_n$, with careful calibration via the use of thresholds. This finally helps us enforce the best response condition of Nash equilibria on (x, y) in (A, B) by comparing total expenditures of buyers $\mathbb{C}(p, s)$ and using implications from such comparisons.

We now formally define the SPP game G in the next section, and then the following sections show the hardness result based on G .

3.1.1. The SPP Game.

The game G has the following set of goods:

- Normalization goods: n goods $\{N(p, s)_1, \dots, N(p, s)_n\}$ for each $p \in \{1, 2\}$ and $s \in [n]$.
- Expenditure goods: n goods $\{E(p, s)_1, \dots, E(p, s)_n\}$ for each $p \in \{1, 2\}$ and $s \in [n]$.
- Threshold goods: 1 good $T(p, s)$ for each $p \in \{1, 2\}$ and $s \in [n]$.

Set $v = 1/(16n)$. The set of buyers in G is defined as follows, where we write $V(\cdot, \cdot)$ to denote the value of a good (the second component) to a buyer (the first component):

- Buyer $\mathbb{C}(p, s)$, $p \in \{1, 2\}$ and $s \in [n]$: $\mathbb{C}(p, s)$ has positive values for the following goods:
 - Normalization goods: $V(\mathbb{C}(p, s), N(p, s)_i) = 2$ for all $i \in [n] \setminus \{s\}$;
 - $V(\mathbb{C}(p, s), N(p, s)_s) = 1$; and $V(\mathbb{C}(p, s), N(p, t)_s) = 1$ for all $t \in [n] \setminus \{s\}$.
 - Threshold good $T(p, s)$: $V(\mathbb{C}(p, s), T(p, s)) = 2n^4$.
 - Expenditure goods: $V(\mathbb{C}(p, s), E(p, s)_i) = 1$ for all $i \in [n]$.

For $p = 1$: $V(\mathbb{C}(1, s), E(2, t)_s) = vB_{st}$ for all $t \in [n]$.

For $p = 2$: $V(\mathbb{C}(2, s), E(1, t)_s) = vA_{ts}$ for all $t \in [n]$.

For $p = 1$, the budget of $\mathbb{C}(1, s)$ is $n/2 + n^4 + 1/4 - v + \sum_{t \in [n]} vA_{st}/2$;

For $p = 2$, the budget of $\mathbb{C}(2, s)$ is $n/2 + n^4 + 1/4 - v + \sum_{t \in [n]} vB_{ts}/2$.

- Threshold Buyer \mathbb{T} : \mathbb{T} has positive values only for the following goods:

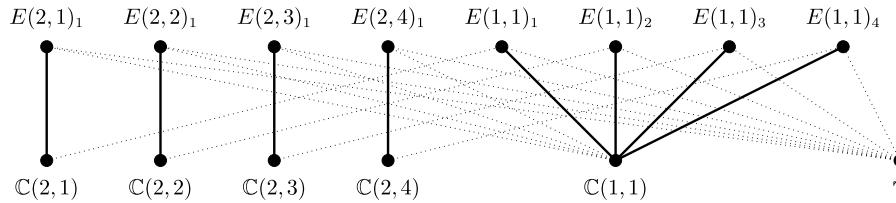
- Threshold goods: $V(\mathbb{T}, T(p, s)) = n^4$ for each $p \in \{1, 2\}$ and $s \in [n]$.
- Expenditure goods: $V(\mathbb{T}, E(1, s)_t) = vA_{st}/2$ and $V(\mathbb{T}, E(2, s)_t) = vB_{ts}/2$ for all $s, t \in [n]$.

\mathbb{T} has budget n^7 (high enough so that $\alpha(\mathbb{T}) = 1$ in any PE).

- Dummy buyers $\mathbb{D}(p, s)$, $p \in \{1, 2\}$ and $s \in [n]$: The budget of $\mathbb{D}(p, s)$ is v and she only values the normalization good $N(p, s)_s$ at $V(\mathbb{D}(p, s), N(p, s)_s) = 1$.

See Figure 1 and Figure 2 for a pictorial representation of the valuation structure. It is clear from the definition of G that it can be constructed from (A, B) in polynomial time.

3.1.2. Structure of Pacing Equilibria of G . With the definition of G in place, we start by showing some auxiliary structural results on the PE of G ; these will be used to construct strategies for the bimatrix game. Let \mathcal{E} be a PE of the SPP game G . We will use $\alpha(b)$ to denote the pacing multiplier of buyer b in \mathcal{E} . Observe that, from the definition of pacing equilibria, we can conclude that $\alpha(\mathbb{T}) = 1$ in \mathcal{E} ; otherwise \mathbb{T} needs to spend all her budget of n^7 , which is

Figure 2. (Color online) All the expenditure goods for which buyer $\mathbb{C}(1,1)$ has a nonzero value, when $n = 4$.

Notes. A buyer having a nonzero value for a good is represented by a line connecting the two. Solid lines denote a value of 1 and dotted lines denote values which are smaller than $v = 1/(16n)$.

impossible given that no buyer has value more than $2n^4$ for any good. The following lemma establishes bounds on $\alpha(\mathbb{C}(p,s))$ in \mathcal{E} .

Lemma 1. For each $p \in \{1, 2\}$ and $s \in [n]$, $1/2 \leq \alpha(\mathbb{C}(p,s)) < 1$ and $\alpha(\mathbb{D}(p,s)) = \alpha(\mathbb{C}(p,s))$.

Suppose for some $p \in \{1, 2\}$ and $s \in [n]$, we have $\alpha(\mathbb{C}(p,s)) < 1/2$. Then $\mathbb{C}(p,s)$ doesn't win any part of threshold good $T(p,s)$. Observe that she has value at most 2 for every other good. Given that there are only $O(n^2)$ goods in G , she cannot possibly spend all her budget (which is $\Omega(n^4)$). Here, we have used the fact that the payment is smaller than her bid on every item that she wins because of the second-price auction format, which in turn is always smaller than her value. This contradicts the assumption that \mathcal{E} is a PE of G . Therefore, $\alpha(\mathbb{C}(p,s)) \geq 1/2$.

Next we prove $\alpha(\mathbb{D}(p,s)) = \alpha(\mathbb{C}(p,s))$. Suppose $\alpha(\mathbb{D}(p,s)) > \alpha(\mathbb{C}(p,s))$ for some $p \in \{1, 2\}$ and $s \in [n]$. Then, buyer $\mathbb{D}(p,s)$ wins all of good $N(p,s)_s$ at price $\alpha(\mathbb{C}(p,s)) \geq 1/2$ because $\mathbb{D}(p,s)$ and $\mathbb{C}(p,s)$ both value $N(p,s)_s$ at 1, and the rest of the buyers have zero value for it. This violates her budget constraint and leads to a contradiction. Therefore, $\alpha(\mathbb{D}(p,s)) \leq \alpha(\mathbb{C}(p,s))$. Moreover, if $\alpha(\mathbb{D}(p,s)) < \alpha(\mathbb{C}(p,s))$ (which implies $\alpha(\mathbb{D}(p,s)) < 1$) then her expenditure is zero. This violates the no unnecessary pacing condition. Hence, $\alpha(\mathbb{D}(p,s)) = \alpha(\mathbb{C}(p,s))$ must hold. Observe that, in particular, this means that the price of $N(p,s)_s$ is $\alpha(\mathbb{C}(p,s))$.

Finally suppose $\alpha(\mathbb{C}(p,s)) = 1$ for some $p \in \{1, 2\}$, $s \in [n]$. Then she wins the following goods:

- All of normalization goods $N(p,s)_t$ for each $t \neq s$ because $\mathbb{C}(p,s)$ has the higher value for them, and she spends at least $1/2$ on each of them because $\alpha(\mathbb{C}(p,t)) \geq 1/2$ by the first part of the proof.
- Part of normalization good $N(p,s)_s$ by spending at least $1 - v$. This is because $N(p,s)_s$ has price 1, she shares it with $\mathbb{D}(p,s)$, and buyer $\mathbb{D}(p,s)$ only has budget v .
- All of threshold good $T(p,s)$ by spending n^4 because she has the higher value.
- All of expenditure good $E(p,s)_t$, for each $t \in [n]$, by spending at least $vA_{st}/2$ if $p = 1$ and $vB_{ts}/2$ if $p = 2$ because she has the higher value.

Hence, the total expenditure of $\mathbb{C}(p,s)$ is at least $(n-1)/2 + 1 - v + n^4 + \sum_t vA_{st}/2$ if $p = 1$ and at least $(n-1)/2 + 1 - v + n^4 + \sum_t vB_{ts}/2$ if $p = 2$. In both cases, the budget constraint is violated, leading to a contradiction. Therefore, the lemma holds.

The above lemma implies that every $\mathbb{C}(p,s)$ is paced in \mathcal{E} (i.e., $\alpha(\mathbb{C}(p,s)) < 1$), thereby implying that their total expenditures must exactly equal their budgets. Additionally, we have the following corollary which will be used in the proof of Lemma 3.

Corollary 1. For each $p \in \{1, 2\}$ and $s \in [n]$, $\mathbb{C}(p,s)$ spends exactly $\alpha(\mathbb{C}(p,s)) - v$ on $N(p,s)_s$.

Next let $x'_s = \alpha(\mathbb{C}(1,s)) - 1/2$ and $y'_s = \alpha(\mathbb{C}(2,s)) - 1/2$ for each $s \in [n]$. The following lemma will allow us to normalize x' and y' to obtain probability distributions x and y .

Lemma 2. The following inequalities hold: $\sum_{s \in [n]} x'_s > 0$ and $\sum_{s \in [n]} y'_s > 0$.

We show $\sum_s x'_s > 0$. The proof of $\sum_s y'_s > 0$ is completely analogous. Suppose $\sum_s x'_s = 0$. Then, $\alpha(\mathbb{C}(1,s)) = 1/2$ for all $s \in [n]$ because $\alpha(\mathbb{C}(1,s)) = 1/2$ by Lemma 1. We argue below that $\mathbb{C}(1,1)$ violates the no-unnecessary-pacing condition.

To see this, observe $\mathbb{C}(1,1)$ only wins a nonzero fraction of the following goods, and spends:

- At most $1/2$ on each normalization good $N(1,1)_t$, $t \in [n]$, because the highest competing bid is $1/2$ on these goods.
- At most n^4 on the threshold good $T(1,1)$ because that is the highest competing bid.
- At most vA_{1t} on each expenditure good $E(1,1)_t$, $t \in [n]$, because that is the highest possible competing bid.

Hence, the total expenditure of $\mathbb{C}(1,1)$ is at most $n/2 + n^4 + \sum_t vA_{1t}$, which is strictly less than her budget of $n/2 + n^4 + 1/4 - v + \sum_t vA_{1t}/2$, a contradiction. \square

3.1.3. Extracting Bimatrix Game Equilibria from G . Now, we are ready to define the mixed strategies (x, y) for the bimatrix game (A, B) . Set player 1's mixed strategy x to be $x_s = x'_s / \sum_i x'_i$ and player 2's mixed strategy y to be $y_s = y'_s / \sum_i y'_i$. These are valid mixed strategies because of Lemmas 1 and 2. The next lemma shows that (x, y) is indeed a Nash equilibrium of (A, B) .

Lemma 3. (x, y) is a Nash equilibrium for the bimatrix game (A, B) .

Suppose there are $s, s^* \in [n]$ such that $x_s > 0$ but $\sum_t A_{st} y_t > \sum_t A_{s^*t} y_t$ (the proof for y is analogous). Using $x_s > 0$, buyer $\mathbb{C}(1, s)$ spends nonzero amounts on the following goods:

- $\alpha(\mathbb{C}(1, t))$ on the normalization good $N(1, s)_t$ for each $t \neq s$ because $\mathbb{C}(1, s)$ has a bid strictly greater than 1, which is the value and an upper bound on the bid of $\mathbb{C}(1, t)$.
- $\alpha(\mathbb{C}(1, s)) - v$ on the normalization good $N(1, s)_s$ because she shares the good with $\mathbb{D}(1, s)$ who has a budget of v .
- n^4 on the threshold good $T(1, s)$ because her bid is strictly greater than n^4 .
- $\alpha(\mathbb{C}(2, t)) \cdot v A_{st}$ on the expenditure good $E(1, s)_t$ for each $t \in [n]$.

Therefore, the total expenditure of buyer $\mathbb{C}(1, s)$ is given by

$$\begin{aligned} & \sum_{t \in [n]} \alpha(\mathbb{C}(1, t)) + n^4 - v + \sum_{t \in [n]} \alpha(\mathbb{C}(2, t)) \cdot v A_{st} \\ &= \sum_{t \in [n]} \alpha(\mathbb{C}(1, t)) + n^4 - v + \sum_{t \in [n]} v A_{st} / 2 + \sum_{t \in [n]} y_t v A_{st} \end{aligned}$$

Note that the RHS above after replacing s with s^* :

$$\sum_{t \in [n]} \alpha(\mathbb{C}(1, t)) + n^4 - v + \sum_{t \in [n]} v A_{s^*t} / 2 + \sum_{t \in [n]} y_t v A_{s^*t}$$

is an upper bound for the total expenditure of buyer $\mathbb{C}(1, s^*)$ (no matter whether $x_{s^*} > 0$ or not).

As a result, the total expenditure of $\mathbb{C}(1, s)$ minus that of $\mathbb{C}(1, s^*)$ is at least

$$\left(\sum_{t \in [n]} v A_{st} / 2 + \sum_{t \in [n]} y_t v A_{st} \right) - \left(\sum_{t \in [n]} v A_{s^*t} / 2 + \sum_{t \in [n]} y_t v A_{s^*t} \right) > \sum_{t \in [n]} v A_{st} / 2 - \sum_{t \in [n]} v A_{s^*t} / 2$$

using the assumption that $\sum_t A_{st} y_t > \sum_t A_{s^*t} y_t$. On the other hand, the budget of $\mathbb{C}(1, s)$ minus that of $\mathbb{C}(1, s^*)$ is equal to the RHS above. This is a contradiction because both buyers should have their total expenditures equal to their budgets. This finishes the proof of the lemma.

Thus, given a $\{0, 1\}$ -cost bimatrix game (A, B) , we have defined an SPP game G which satisfies the following properties: (i) G can be constructed in polynomial time; (ii) any PE \mathcal{E} of G can be used to construct a Nash equilibrium (x, y) of (A, B) in polynomial time. As a result, the problem of finding an exact pacing equilibrium in a second-price pacing game is PPAD-hard.

3.2. Hardness of Finding Approximate Pacing Equilibria

We next state our main hardness result, which extends the PPAD-hardness of finding pacing equilibria to the approximate case of finding (δ, γ) -approximate pacing equilibria.

Theorem 3. The problem of computing a (δ, γ) -approximate PE of an SPP game $G = (n, m, (v_{ij}), (B_i))$ with $\delta = \gamma = 1/n^7$ is PPAD-hard.

The proof is relegated to Appendix A. It uses similar ideas but entails more involved bookkeeping to incorporate approximations introduced in (δ, γ) -approximate PE. Theorem 2 follows from Theorem 3 by standard padding arguments (i.e., adding dummy buyers to the game).

3.3. Implications of the Hardness Result

Before concluding this section, we discuss some implications of our hardness results. In Borgs et al. [9], the authors introduced a natural bidding heuristic for optimizing the utility of budget-constrained agents who repeatedly participate in day-long auction campaigns for m items, where the set of agents and items remains the same every day. The heuristic maintains a pacing multiplier for each agent, which is increased by a small amount if the buyer ran out of her daily budget before the end of the previous day, and decreased otherwise. They use random perturbation to avoid instabilities, which gives an agent who bids close to the highest bid a fraction of the item in expectation. If we ignore the intraday temporal aspects of their model, their setting can be thought of as repeatedly playing the

perturbed SPP game from Section 2.1 every day. In Theorem 1 of Borgs et al. [9], they prove that their heuristic efficiently converges for first-price auctions. Furthermore, they conjecture the convergence of the heuristic for second-price auctions to pacing multipliers which satisfy the following conditions: (i) Every agent runs out of her daily-budget close to the end of the day; (ii) Every agent either spends most of her daily budget or has a pacing multiplier close to one. In Theorem D.1 of Appendix D, we show that Theorem 3 implies that computing an approximate pacing equilibrium of the perturbed SPP game is also PPAD hard. As a consequence, if $\text{PPAD} \neq \text{P}$, then ALGORITHM 1 of Borgs et al. [9] does not always converge efficiently for second-price auctions, that is, the number of days/time-steps required for convergence cannot scale as a polynomial function of the input size and $(1/\delta, 1/\gamma)$ in the worst-case. In other words, we have shown that theorem 1 of Borgs et al. [9] cannot be extended to second-price auctions in any way that maintains efficient convergence unless $\text{PPAD} = \text{P}$, thereby making progress toward their open conjecture.

Moreover, recall from Section 2.1 that if all of the buyers employ pacing algorithms, like the one proposed by Balseiro and Gur [4], and the resulting dynamics converge, then they will converge to an approximate pacing equilibrium. Our hardness result (Theorems 3 and D.1) implies that there exists a (correlated) value distribution such that the algorithm of Balseiro and Gur [4], which is optimal for a single buyer against an adversarial/stochastic competition, does not converge efficiently to an equilibrium when employed by all the buyers, unless $\text{PPAD} = \text{P}$.

Our hardness results are also pertinent to the relationship between pacing equilibria and market equilibria. In proposition 5 of Conitzer et al. [18], the authors show that every pacing equilibrium in a second-price pacing game has an equivalent supply-aware market equilibrium with linear utilities, where supply-aware means that the buyers are aware of the supplies of each item and choose their demand set accordingly. Thus, the relationship between pacing equilibria and market equilibria, in combination with Theorem 3, implies that there exists a refinement of the set of supply-aware market equilibria with linear utilities which is PPAD-hard to compute.

4. Existence of Pacing Equilibria and Membership in PPAD

We prove Theorem 1 in this section, that is, the problem of finding a pacing equilibrium of an SPP game is in PPAD. One consequence of this result is that every SPP game with rational values v_{ij} and budgets B_i has a pacing equilibrium (α, x) with rational entries.

Our plan is as follows. We first introduce a restricted version of approximate pacing equilibria called *smooth* (δ, γ) -approximate PE (see Definition 3), which will only be used in Section 4.1. We prove in Section 4.1 that the problem of finding a smooth (δ, γ) -approximate PE (when δ and γ are input parameters encoded in binary) is in PPAD. Given that the smooth version (Definition 3) is a restriction of (δ, γ) -approximate PE (Definition 2), this implies that the problem of computing a (δ, γ) -approximate PE is in PPAD.

Next we give in Section 4.2 an efficient algorithm that can round any $(\delta, \gamma/2)$ -approximate PE into a γ -approximate PE when δ is sufficiently small. This, combined with the PPAD-membership of (δ, γ) -approximate PE, shows that the problem of computing γ -approximate PE is also in PPAD.

Finally we show in Section 4.3 that, when γ is sufficiently small, any γ -approximate PE of G can be used to build a linear program which can then be solved to obtain an *exact* pacing equilibrium of G . It follows that the problem of computing an exact pacing equilibrium is in PPAD.

4.1. PPAD Membership of Computing (δ, γ) -Approximate Equilibria

We start with the definition of *smooth* (δ, γ) -approximate PE. It is a refinement of (δ, γ) -approximate PE in which the pacing multipliers (α_i) *fully determine* the allocations (x_{ij}) . Note that this is not the case for (δ, γ) -approximate PE in general: potentially there can be (δ, γ) -approximate PE with identical multipliers but different allocations. The smooth version we consider below, on the other hand, specifies the allocations as continuous functions of multipliers.

Definition 3 (Smooth Approximate Pacing Equilibria). Given an SPP game $G = (n, m, (v_{ij}), (B_i))$ and two parameters $\delta \in (0, 1), \gamma \in [0, 1)$, we say that (α, x) with $\alpha = (\alpha_i) \in [0, 1]^n$, $x = (x_{ij}) \in [0, 1]^{nm}$ and $\sum_{i \in [n]} x_{ij} \leq 1$ for all $j \in [m]$ is a *smooth* (δ, γ) -approximate PE of G if

(a) Only buyers close to the highest bid win the good and the allocation x is completely specified by α : For each $i \in [n]$ and $j \in [m]$, x_{ij} (as a function of α) is given by

$$x_{ij}(\alpha) := \frac{[\alpha_i v_{ij} - (1 - \delta)h_j(\alpha)]^+}{\sum_{r \in [n]} [\alpha_r v_{rj} - (1 - \delta)h_j(\alpha)]^+}$$

where $[y]^+$ is y if $y \geq 0$ and 0 otherwise. (We assume by default that $0/0 = 0$.)

(b) Budgets are satisfied: $\sum_{j \in [m]} x_{ij}(\alpha) p_j(\alpha) \leq B_i$.

(c) Not too much unnecessary pacing: $\sum_{j \in [m]} x_{ij}(\alpha)p_j(\alpha) < (1 - \gamma)B_i$ implies $\alpha_i \geq 1 - \gamma$.

Observe from the definition that, if (α, x) is a smooth (δ, γ) -approximate PE of an SPP game G , then it must be a (δ, γ) -approximate PE of G as well. Therefore, the PPAD membership of computing a smooth (δ, γ) -approximate PE in an SPP game implies directly the PPAD membership for (δ, γ) -approximate PE. A similar statement holds for establishing their existence.

The main tools we will use are Sperner's Lemma and the search problem it defines.

4.1.1. High-Dimensional Sperner's Lemma. We review Sperner's lemma. Consider a $(n - 1)$ -dimensional simplex $S = \{\sum_{i=1}^n \alpha_i v_i \mid \alpha_i \geq 0, \sum_{i=1}^n \alpha_i = 1\}$, where v_1, \dots, v_n are n vertices of S . A triangulation of S is a partition of S into smaller subsimplices such that any two subsimplices either are disjoint or share a full face of a certain dimension. A *Sperner coloring* T of a triangulation of S is then an assignment of n colors $\{1, \dots, n\}$ to vertices of the triangulation (union of the vertices of subsimplices that make up the triangulation) such that

- Vertices of the original simplex S each receive a different color: $T(v_i) = i$ for each $i \in [n]$.
- Vertices on each face of S are colored using only the colors of the vertices defining that face: For any vertex $u = \sum_i \beta_i v_i$ in the triangulation, we have $T(u) \neq j$ if $\beta_j = 0$.

A *panchromatic* subsimplex of T is one in the triangulation whose vertices have all the n colors.

4.1.2. Sperner's Lemma. Every Sperner coloring T of any triangulation of S has a panchromatic subsimplex.

Before proceeding with the formal proof of PPAD membership (with its added burden of rigorously attending to complexity-theoretic details), we provide an informal argument for the existence of smooth (δ, γ) -approximate PE which forms the basis of its PPAD membership proof. Let G be an SPP game and S be the standard simplex $S = \{\beta = (\beta_1, \dots, \beta_n) \mid \beta_i \geq 0, \sum_i \beta_i = 1\}$ from now on. We will assign a color to each point $\beta \in S$ (informally) as follows: Construct a pacing multiplier $\alpha_i(t) = t\beta_i$ for each $i \in [n]$, where t is a scalar. Increase t , starting at 0, and instruct each buyer $i \in [n]$ to say "Stop" when either $\alpha_i(t) = 1$ or $\sum_j x_{ij}(\alpha(t))p_j(\alpha(t)) = B_i$ happens. Color β with k if buyer k is the first to say "Stop" (with tie breaking done arbitrarily, e.g., taking the smallest such k).

Let $t^*(\beta)$ be the value of t at which some buyer says "Stop" for the first time. Then the buyer that says "Stop" first is either spending her budget or is not paced, that is, she satisfies both the budget constraint (b) and the 'No unnecessary pacing' condition (c) (see Definition 1). Now, by taking a triangulation of S , it is easy to verify that the coloring described above induces a Sperner coloring and thus, Sperner's lemma implies the existence of a panchromatic subsimplex Q . It follows from our coloring that every buyer says "Stop" at one of the vertices of Q and hence, every buyer satisfies (b) and (c) of Definition 1 at one of its vertices. By proving the Lipschitzness of $t^*(\beta)$ and the total expenditures of buyers, both as functions of β , we show that when the triangulation is fine enough, any point β in a panchromatic subsimplex yields a (δ, γ) -approximate PE of G .

With the blueprint of the proof in place, we now proceed with the formal proof that places the problem of computing smooth (δ, γ) -approximate PE in PPAD. Let S be the standard simplex as above, and we consider *Kuhn's triangulation* of S (Deng et al. [21], Kuhn [33]). Given any $\omega > 0$ with $1/\omega$ being an integer, Kuhn's triangulation uses S_ω as its vertices, where S_ω consists of all points $\beta \in S$ whose coordinates β_i are integer multiples of ω . Kuhn's triangulation also has the property that any two vertices of a subsimplex of the triangulation has ℓ_∞ -distance at most 2ω .

A proof of the following PPAD membership result can be found in Etessami and Yannakakis [24] (see the proof of item 2 of Proposition 2.2; note that on p. 2548 they reduce the problem they are interested in to the problem of finding a panchromatic subsimplex in a Sperner coloring over Kuhn's triangulation and then show the latter is in PPAD):

Theorem 4. *Given a Boolean circuit⁵ that encodes a Sperner coloring $T : S_\omega \rightarrow [n]$ of Kuhn's triangulation for some ω and n , the problem of finding a panchromatic subsimplex is in PPAD.*

We prove the PPAD membership of the problem of finding a smooth (δ, γ) -approximate PE by giving a polynomial-time reduction to the problem described in Theorem 4. Given an SPP game $G = (n, m, (v_{ij}), (B_i))$ and parameters δ and γ (which we assume without loss of generality that $\delta, \gamma < 1/4$), we set the parameter ω to be

$$\omega = \frac{\min(B_{\min}, 1)}{(2^{|G|}/\delta)^{10,000}} \cdot \frac{\gamma}{2}$$

where $B_{\min} := \min_{i \in [n]} B_i$ and $|G|$ denotes the number of bits needed to represent G . We define a coloring $T : S_\omega \rightarrow [n]$, following ideas described in the sketch of existence above, and prove that T satisfies the following properties:

Lemma 4.

1. T is a Sperner coloring;
2. Every panchromatic subsimplex of T in the triangulation can be used to compute a smooth (δ, γ) -approximate PE of the SPP game G in polynomial time.
3. There is a polynomial-time algorithm that outputs $T(\beta)$ on inputs G, ω, δ and $\beta \in S_\omega$.

The PPAD membership of computing a smooth (δ, γ) -approximate PE in an SPP game follows directly by combining Theorem 4 and Lemma 4.

We now give the definition of the coloring $T : S_\omega \rightarrow [n]$. Let $\beta = (\beta_1, \dots, \beta_n)$ be a vertex of S_ω . Set $\alpha_i(t) = t\beta_i$, where t is a positive scalar. As discussed earlier, we set the color $T(\beta)$ of β by increasing t , starting at 0, and instructing each buyer i to say “Stop” when either $\alpha_i(t) = 1$ or

$$\sum_{j \in [m]} x_{ij}(\alpha(t)) \cdot p_j(\alpha(t)) = B_i.$$

The color $T(\beta)$ of β is set to be $k \in [n]$ if buyer k is the first buyer to say “Stop” (with arbitrary tie breaking, e.g., by taking the smallest such k).

More formally, recall that for $t > 0$,

$$x_{ij}(\alpha(t)) = \frac{[t\beta_i v_{ij} - (1 - \delta)\max_k t\beta_k v_{kj}]^+}{\sum_r [t\beta_r v_{rj} - (1 - \delta)\max_k t\beta_k v_{kj}]^+} = \frac{[\beta_i v_{ij} - (1 - \delta)\max_k \beta_k v_{kj}]^+}{\sum_r [\beta_r v_{rj} - (1 - \delta)\max_k \beta_k v_{kj}]^+} = x_{ij}(\beta),$$

which does not depend on t . Also, for $t \geq 0$, $p_j(\alpha(t)) = tp_j(\beta)$, where we write $p_j(\beta)$ to denote the second largest element among $\beta_1 v_{1j}, \dots, \beta_n v_{nj}$. For each buyer $i \in [n]$, define

$$t_i(\beta) = \min \left\{ \frac{1}{\beta_i}, \frac{B_i}{\sum_j x_{ij}(\beta) p_j(\beta)} \right\},$$

where the first term is $+\infty$ if $\beta_i = 0$ and the second term is $+\infty$ if $\sum_j x_{ij}(\beta) p_j(\beta) = 0$. Note that $t_i(\beta)$ is exactly the value of t at which buyer i would say “Stop” in the informal coloring procedure described earlier. Given our assumption of $B_i > 0$, we have $t_i(\beta) > 0$ for all $i \in [n]$. Additionally, define $t^*(\beta) = \min_{i \in [n]} t_i(\beta)$. Given that β_i ’s sum to 1, we have that $t^*(\beta) \leq n$ because $\beta_i \geq 1/n$ for some $i \in [n]$. We record the discussion as the following lemma:

Lemma 5. For every $\beta \in S_\omega$ we have $0 < t^*(\beta) \leq n$.

Finally, the color $T(\beta)$ of $\beta \in S_\omega$ is set to be the smallest $i \in [n]$ such that $t_i(\beta) = t^*(\beta)$. We are now ready to prove Lemma 4.

Proof of Lemma 4. Part (3) of Lemma 4 follows from the description of T . To prove part (1) (T is a Sperner coloring), consider a vertex $\beta \in S_\omega$ on the facet of S opposite to the vertex e_i , that is, $\beta_i = 0$. Hence, $t_i(\beta) = \infty$, which by Lemma 5 implies that $T(\beta) \neq i$ given that $t^*(\beta) \leq n$.

To prove part (2), we show that if q is a vertex of any panchromatic subsimplex of T , then (α, x) must be a smooth (δ, γ) -approximate PE of G where $\alpha = t^*(q) \cdot q$ and $x = (x_{ij})$ has $x_{ij} = x_{ij}(q)$.

First it follows from the definition of $t^*(\beta)$ and $x_{ij}(\beta)$ that $\alpha_i \in [0, 1]$ and $x_{ij} \in [0, 1]$. Conditions (a) and (b) of Definition 3 also trivially hold for all vertices of the triangulation. It suffices to prove (c) for all $i \in [n]$, which means the complementarity condition that either $\alpha_i \geq 1 - \gamma$ or the expenditure of buyer i is at least $(1 - \gamma)B_i$. Fix an arbitrary $i \in [n]$.

For this purpose we note that given the subsimplex is panchromatic, it has a vertex q' such that $T(q') = i$, which implies that if we used q' to define α' and x' (i.e., $\alpha' = t^*(q') \cdot q'$ and $x'_{ij} = x_{ij}(q')$), then they would satisfy the above complementarity condition for buyer i with $\gamma = 0$. The following claim shows that both the multiplier $t^*(\beta) \cdot \beta_i$ and the total expenditure of buyer i :

$$\sum_{j \in [m]} x_{ij}(\beta) \cdot p_j(t^*(\beta) \cdot \beta) = t^*(\beta) \sum_{j \in [m]} x_{ij}(\beta) \cdot p_j(\beta)$$

are smooth as functions of β . Intuitively this allows us to use the complementarity condition for buyer i at q' to show that the same condition holds at q approximately given that $\|q - q'\|_\infty \leq 2\omega$ (as a property of subsimplices in Kuhn’s triangulation).

Claim 1. Let $L = (2^{|G|}/\delta)^{10,000}$. Then for any panchromatic subsimplex S_0 of T and buyer $i \in [n]$, the following Lipschitz conditions hold for all $\beta, \beta' \in S_0$:

$$\begin{aligned} |t^*(\beta) \cdot \beta_i - t^*(\beta') \cdot \beta'_i| &\leq L \cdot \|\beta - \beta'\|_\infty \quad \text{and} \\ \left| t^*(\beta) \sum_{j \in [m]} x_{ij}(\beta) \cdot p_j(\beta) - t^*(\beta') \sum_{j \in [m]} x_{ij}(\beta') \cdot p_j(\beta') \right| &\leq L \cdot \|\beta - \beta'\|_\infty \end{aligned}$$

We use Claim 1 to finish the proof of the lemma and consign the claim's proof to Appendix B. Given $T(q') = i$, one of the following two cases holds:

- $t^*(q') \cdot q'_i = 1$, which by Claim 1 and our choice of ω implies

$$\alpha_i = t^*(q) \cdot q_i \geq 1 - 2L\omega \geq 1 - \gamma$$

- $t^*(q') \sum_j x_{ij}(q') p_j(q') = B_i$, which in combination with Claim 1 and our choice of ω implies that the expenditure of buyer i exceeds $(1 - \gamma)B_i$:

$$t^*(q) \sum_{j \in [m]} x_{ij}(q) \cdot p_j(q) \geq B_i - 2L\omega \geq B_i - B_{\min} \gamma \geq (1 - \gamma)B_i.$$

Since $i \in [n]$ was arbitrary, this finishes the proof that (α, x) is a smooth (δ, γ) -approximate PE. \square

4.2. PPAD Membership of Computing γ -Approximate PE

Consider an SPP game $G = (n, m, \{v_{ij}\}_{i,j}, \{B_i\}_i)$. As before, we will use $|G|$ to denote the number of bits required to represent G . The main result of this subsection shows that (informally) when δ is small enough, any $(\delta, \gamma/2)$ -approximate PE of G can be efficiently rounded to a γ -approximate PE. It follows from the PPAD membership of (δ, γ) -approximate PE established in the previous subsection that the problem of computing a γ -approximate PE is in PPAD as well.

Before presenting the rounding algorithm, we motivate the main idea behind it. Observe that the major difference between (δ, γ) -approximate PE and γ -approximate PE is the ability of buyers that don't have the highest bid to win the good in the former. In order to round a (δ, γ') -approximate PE (α^*, x^*) to obtain a γ -approximate PE (α', x') of G (where $\gamma' = \gamma/2$ in the rest of this subsection), we set $x' = x^*$ and need to round α^* to α' to ensure that all the winners are tied for the highest bid and at the same time, the multiplier and total expenditure of each buyer changes only slightly.

We now present an informal argument that demonstrates how this is achieved in our rounding algorithm when there are only two buyers ($n = 2$). Define the set of all valuation ratios

$$\tilde{\mathcal{V}} = \left\{ \frac{v_{ar}}{v_{br}} : a, b \in [n], r \in [m] \text{ such that } v_{ar}, v_{br} > 0 \right\}.$$

Set δ to be small enough: for all $y, z \in \mathcal{V}$ with $yz > 1$, we have $(1 - \delta)^2 yz > 1$. Consider a (δ, γ') -approximate PE (α^*, x^*) . Assume without loss of generality that there is a good j such that $\alpha_1^* v_{1j} = c_j \alpha_2^* v_{2j}$ and $1 - \delta \leq c_j \leq 1/(1 - \delta)$. If no such good j exists then every good is fully allocated to the buyer with the highest bid because only bidders with bids greater than $(1 - \delta)$ times the highest bid can win the item in a (δ, γ') -approximate PE, and thus, (α^*, x^*) is already a γ' -approximate PE. We show that after scaling the pacing multiplier of buyer 2 from α_2^* to $c_j \alpha_2^*$ (and letting $\alpha' = (\alpha_1^*, c_j \alpha_2^*)$ be the new multipliers), (α', x^*) satisfies the property that $x_{i\ell}^* > 0$ for any i and ℓ implies buyer i has the highest bid for good ℓ . This is trivially true for good $\ell = j$ given that the two buyers are now tied on good j . The remaining goods can be divided into two categories and we argue about each one separately:

- Consider good ℓ such that $\alpha_1^* v_{1\ell} = c_\ell \alpha_2^* v_{2\ell}$ and c_ℓ satisfies either $c_\ell < 1 - \delta$ or $c_\ell > 1/(1 - \delta)$. Given that we only changed the multiplier of buyer 2 by a factor of $1 - \delta \leq c_j \leq 1/(1 - \delta)$, the highest bidder does not change. Moreover, the highest bidder won the entire good in the (δ, γ') -approximate PE because $1 - \delta \leq c_j \leq 1/(1 - \delta)$ and continues to do so in the (δ, γ') -approximate PE because the allocation does not change.

- Consider a good ℓ such that $\alpha_1^* v_{1\ell} = c_\ell \alpha_2^* v_{2\ell}$ and c_ℓ satisfies $(1 - \delta) \leq c_\ell \leq 1/(1 - \delta)$. Then, we can write $\alpha_1^*/\alpha_2^* = c_j(v_{2j}/v_{1j}) = c_\ell(v_{2\ell}/v_{1\ell})$, which implies $(c_j/c_\ell)(v_{2j}/v_{1j})(v_{1\ell}/v_{2\ell}) = 1$. Observe that $c_j/c_\ell \in [(1 - \delta)^2, 1/(1 - \delta)^2]$. Thus, by our choice of δ , we get $(v_{2j}/v_{1j})(v_{1\ell}/v_{2\ell}) = 1$, which implies $c_j = c_\ell$. Hence, both buyers are tied in good ℓ .

To finish the proof that (α', x^*) is a γ -approximate PE, it suffices to show that the budget constraint and the not too much unnecessary pacing condition still hold approximately after the small scaling of α_2^* . In the rest of this subsection, we extend the aforementioned line of reasoning to design a rounding algorithm for the general setting, and prove its correctness.

Building on $\tilde{\mathcal{V}}$ defined above, we can define the set of valuation ratio products

$$\mathcal{V} = \{y_1 y_2 \dots y_k : k \in [2n] \text{ and } y_i \in \tilde{\mathcal{V}} \text{ for each } i \in [k]\},$$

that is, \mathcal{V} consists of all products of no more than $2n$ numbers from $\tilde{\mathcal{V}}$. Given G and $\gamma \in [0, 1)$ (with $\gamma' = \gamma/2$), we choose $\delta \in [0, 1)$ to be small enough to satisfy the following two conditions:

$$(1 - \delta)^{2^n} > (1 - \gamma') \text{ and } (1 - \delta)^{2^n} z > 1 \text{ for all } z \in \mathcal{V} \text{ such that } z > 1.$$

Algorithm 1 (Rounding Algorithm)

Initialize: Graph $\mathcal{G} = (V, E)$ with $V = [n]$ and $E = \emptyset$; $\alpha = \alpha^*$

While there exists a good $j \in [m]$ and a buyer $i \in W_j$ such that $\alpha_i v_{ij} < h_j(\alpha)$, that is, i does not have the highest bid on j but wins a positive fraction of it:

1. Pick $k, i \in [n]$ and $j \in [m]$ such that $i \in W_j$ and $\alpha_i v_{ij} < \alpha_k v_{kj} = h_j(\alpha)$
2. Set $\alpha_a \leftarrow (h_j(\alpha)/\alpha_i v_{ij}) \cdot \alpha_a$ for every buyer $a \in C_{\mathcal{G}}(i)$
3. Set $E \leftarrow E \cup \{\{i, k\}\}$ and $I(\{i, k\}) = j$

Return: $\alpha' := (1 - \delta)^{2^n} \alpha$

It suffices to set δ to be $1/2^N$ where N is polynomial in $|G|$ and $\log(1/\gamma)$.

Let (α^*, x^*) be a (δ, γ') -approximate PE of $G = (n, m, (v_{ij}), (B_i))$, where $\gamma' = \gamma/2$ and δ satisfies the two conditions above. We will use W_j to denote the winners of the good j under x^* : W_j consists of buyers i with $x_{ij}^* > 0$. Moreover, recall that $h_j(\alpha)$ denotes the highest bid on good j when the pacing multipliers are given by α . Our rounding algorithm is presented in Algorithm 1. The polynomial reduction then follows from the following performance guarantee of the rounding algorithm, which we prove in the rest of the subsection:

Lemma 6 (Correctness). *The rounding algorithm takes (α^*, x^*) , δ and G as input and runs in polynomial time. Let α' be the tuple of multipliers returned by the rounding algorithm. Then (α', x') is a γ -approximate PE of G .*

The rounding algorithm maintains an undirected graph \mathcal{G} over vertices $[n]$ as buyers. \mathcal{G} starting out with an empty edge set and edges are added according to Algorithm 1 to keep track of the rounding-updates performed on α . We use $C_{\mathcal{G}}(i)$ to denote the connected component of i in the graph \mathcal{G} . The algorithm also maintains an edge labeling $I(\cdot)$ that maps each edge of the graph \mathcal{G} to a good $j \in [m]$ (which intuitively is the good that caused the creation of this edge). We remark that the labeling $I(\cdot)$ is only relevant for the analysis of the algorithm below. Now, we proceed to prove Lemma 6.

Lemma 7. *Suppose in the t_0 iteration of the while loop, $\{i, k\}$ is the edge that was just added to \mathcal{G} with $I(\{i, k\}) = j$, then at the end of this iteration we have $C_{\mathcal{G}}(i) = C_{\mathcal{G}}(k)$ and*

$$\frac{\alpha_i}{\alpha_k} = \frac{v_{kj}}{v_{ij}}. \quad (\#)$$

Moreover, (#) holds for all iterations $t \geq t_0$.

We prove the lemma using induction on the iterations on the while loop. For the base case $t = t_0$, note that (#) holds at the end of the iteration due to Step 2 of Algorithm 1. Moreover, since edge $\{i, k\}$ is added to \mathcal{G} in Step 3, we also have $C_{\mathcal{G}}(i) = C_{\mathcal{G}}(k)$ at the end of iteration t_0 . Moreover, since no edges are removed during the run of Algorithm 1, $\{i, k\} \in E$ for iterations after t_0 , and hence $C_{\mathcal{G}}(i) = C_{\mathcal{G}}(k)$ at the end of all iterations $t \geq t_0$. Suppose (#) holds at the end of iteration $t - 1$ for some $t - 1 \geq t_0$. Then, either both α_i and α_k will both be updated identically or neither of them will be updated because $C_{\mathcal{G}}(i) = C_{\mathcal{G}}(k)$, thereby maintaining (#). This completes the induction and establishes the lemma.

Next we prove that at the end of each iteration, bids for the same good from buyers in the same component of \mathcal{G} are either tied or not very close.

Lemma 8. *After each iteration of the while loop, and for each good $j \in [m]$, all buyers from the same connected component of \mathcal{G} are either tied for j , or their bids for j are multiplicatively separated by a factor larger than $(1 - \delta)^{2^n}$.*

Let \mathcal{G} be the current graph and $a, b \in [n]$ be two buyers in the same connected component of \mathcal{G} . Assuming $\alpha_a v_{aj} > \alpha_b v_{bj}$ for some $j \in [m]$, we show below that $(1 - \delta)^{2^n} \alpha_a v_{aj} > \alpha_b v_{bj}$ from which the lemma follows. Given that a and b are connected in \mathcal{G} , we write $\{a, i_1\}, \{i_1, i_2\}, \dots, \{i_L, b\}$ to denote a path from a to b in \mathcal{G} with $L < n$. Then, using Lemma 7, we can write

$$1 > \frac{\alpha_b v_{bj}}{\alpha_a v_{aj}} = \frac{v_{bj}}{v_{aj}} \cdot \frac{\alpha_{i_1}}{\alpha_a} \frac{\alpha_{i_2}}{\alpha_{i_1}} \frac{\alpha_{i_3}}{\alpha_{i_2}} \dots \frac{\alpha_{i_L}}{\alpha_{i_{L-1}}} = \frac{v_{bj}}{v_{aj}} \cdot \frac{v_{aI(\{a, i_1\})}}{v_{i_1I(\{a, i_1\})}} \frac{v_{i_1I(\{i_1, i_2\})}}{v_{i_2I(\{i_1, i_2\})}} \dots \frac{v_{i_LI(\{i_L, b\})}}{v_{bI(\{i_L, b\})}}$$

Hence, $\alpha_a v_{aj} / \alpha_a v_{bj} \in \mathcal{V}$ and $\alpha_a v_{aj} / \alpha_b v_{bj} > 1$. Therefore, our choice of δ implies that

$$(1 - \delta)^{2^n} \cdot \frac{\alpha_a v_{aj}}{\alpha_b v_{bj}} > 1,$$

as required. \square

Initially (in α^*) we have every $i \in W_j$ has $\alpha_i^* v_{ij} \geq (1 - \delta)h_j(\alpha^*)$ (given that (α^*, x^*) is a (δ, γ') -approximate PE). The next lemma shows that, at the end of each iteration, $\alpha_i v_{ij}$ of every $i \in W_j$ (note that W_j is always defined using the original allocation x^*) remains not far from $h_j(\alpha)$.

Lemma 9. *After t iterations of the while loop, every $j \in [m]$ and $i \in W_j$ satisfy*

$$\alpha_i v_{ij} \geq (1 - \delta)^{2^t} \cdot h_j(\alpha).$$

The proof follows from induction. The base case of $t = 0$ follows from definition.

Suppose the statement holds after $(t - 1)$ iterations, and let's focus on some $j \in [m]$ and $i \in W_j$ during the t -th iteration. By our inductive hypothesis, we have

$$\alpha_i v_{ij} \geq (1 - \delta)^{2^{t-1}} \cdot h_j(\alpha)$$

before the start of the t -th iteration. On the other hand, note that all changes to α occur in step 2 of the while loop, and moreover, all such changes result in an increase of some entries of α . It also follows from the inductive hypothesis and the choices of k, i, j in step 1 of the while loop that entries of α can only go up by a multiplicative factor of at most $1/(1 - \delta)^{2^{t-1}}$. Therefore, after the t -th iteration, we have

$$\alpha_i v_{ij} \geq (1 - \delta)^{2^{t-1}} \cdot (1 - \delta)^{2^{t-1}} \cdot h_j(\alpha) = (1 - \delta)^{2^t} \cdot h_j(\alpha).$$

This completes the induction step. \square

Lemmas 8 and 9 imply that, in each of the first n iterations of the while loop, buyers i and k picked in step 1 must belong to different connected components of \mathcal{G} . As a result, there are at most $n - 1$ iterations of the while loop given that we merge two connected components in each loop. On the one hand, this implies that the rounding algorithm terminates in polynomial time. On the other hand, at the termination of the while loop, for every good $j \in [m]$, we have $\alpha_i v_{ij} = h_j(\alpha)$ for all $i \in W_j$, that is, every winner of j under x^* has the highest bid for j .

The next lemma shows that the α' returned by the rounding algorithm is close to α^* .

Lemma 10. *Let α' be the tuple of multipliers returned by the rounding algorithm. Then*

$$(1 - \delta)^{2^n} \alpha^* \leq \alpha' \leq \alpha^*.$$

By Lemma 9, in iteration t of the while loop, each entry of α either stays the same or increases multiplicatively by a factor of at most $1/(1 - \delta)^{2^{t-1}}$. As there are at most $n - 1$ iterations of the while loop, we have for every $i \in [n]$:

$$(1 - \delta)^{2^n} \cdot \alpha_i^* \leq \alpha'_i := (1 - \delta)^{2^n} \cdot \alpha_i \leq (1 - \delta)^{2^n} \prod_{t=1}^{n-1} \frac{1}{(1 - \delta)^{2^{t-1}}} \cdot \alpha_i^* \leq \alpha_i^*.$$

This finishes the proof of the lemma.

We are now ready to prove Lemma 6.

Proof of Lemma 6. We have already shown that the algorithm runs in polynomial time. Assuming that (α^*, x^*) is a (δ, γ') -approximate PE of \mathcal{G} , we show that (α', x^*) is a γ -approximate PE of \mathcal{G} by establishing conditions (a)-(d) of Definition 2. Using Lemma 10, we have $\alpha' \in [0, 1]^n$. Condition (a) has already been established earlier using Lemmas 8 and 9. Condition (b) holds because we kept the same allocation x^* and given how we obtain α' from α^* , the set of goods j with $h_j(\alpha^*) > 0$ is the same as that in α' . Condition (c) follows easily from Lemma 10. So it suffices to verify that (d) holds with γ .

To see this we have for each buyer $i \in [n]$ that either $\alpha_i^* \geq 1 - \gamma'$ or $\sum_j x_{ij}^* p_j(\alpha^*) \geq (1 - \gamma')B_i$. For the former case, we have from Lemma 10 that

$$\alpha'_i \geq (1 - \delta)^{2^n} \cdot (1 - \gamma') > (1 - \gamma')^2 \geq 1 - \gamma$$

using $(1 - \delta)^{2^n} > 1 - \gamma'$ from the choice of δ and that $\gamma = 2\gamma'$. For the latter case, it follows from Lemma 10 and our

choice of δ that

$$p_j(\alpha') \geq (1 - \delta)^{2^n} \cdot p_j(\alpha^*) > (1 - \gamma') \cdot p_j(\alpha^*)$$

for all $j \in [m]$. Here we have used the fact that $p_j((1 - \delta)^{2^n} \alpha^*) = (1 - \delta)^{2^n} p_j(\alpha^*)$ and $p_j(\alpha) \geq p_j(\tilde{\alpha})$ whenever $\alpha \geq \tilde{\alpha}$. As a result, the total expenditure of buyer i in (α', x^*) is

$$\sum_{j \in [m]} x_{ij}^* \cdot p_j(\alpha') > (1 - \gamma') \sum_{j \in [m]} x_{ij}^* \cdot p_j(\alpha^*) \geq (1 - \gamma')^2 B_i \geq (1 - \gamma) B_i.$$

Therefore, we have shown that (α', x^*) is a γ -approximate PE of G . \square

4.3. PPAD Membership of Computing Exact Pacing Equilibria

In the last subsection we showed that the problem of finding a γ -approximate PE of a second-price pacing game G is in PPAD. Finally we show in this subsection that the problem of finding an exact equilibrium of a pacing game is also in PPAD. To this end, we show that when γ is small enough (though with bit length polynomial in $|G|$), any γ -approximate PE (α', x') of G can be “rounded” into an exact equilibrium by solving a linear program defined using support information extracted from (α', x') . This technique is similar to the one used in Etessami and Yannakakis [24], Vazirani and Yannakakis [41], and Filos-Ratsikas et al. [27]. For this purpose we recall the following fact about linear programs:

Fact 1. There is a polynomial $r(\cdot)$ with the following property. Let LP be a linear program that minimizes a non-negative variable γ . Then an optimal solution of LP has either $\gamma = 0$ or $\gamma \geq 1/2^{r(|\text{LP}|)}$, where $|\text{LP}|$ denotes the number of bits needed to represent LP .

Given a γ -approximate PE (α', x') of $G = (n, m, (v_{ij}), (B_i))$ (for some sufficiently small γ to be specified later), we extract from (α', x') the following support information:

1. $I' \subseteq [n]$ consists of buyers $i \in [n]$ who are almost unpaced, that is, $\alpha'_i \geq 1 - \gamma$. Given that (α', x') is a γ -approximate PE, condition (d) of Definition 2 implies that

$$\sum_{j \in [m]} x'_{ij} p_j(\alpha') \geq (1 - \gamma) B_i, \quad \text{for all } i \notin I.$$

2. For each $j \in [m]$, W'_j is the set of buyers $i \in [n]$ with $x'_{ij} p_j(\alpha') > 0$ (which implies $\alpha'_i v_{ij} = h_j(\alpha')$). These are buyers who win good j and pay a positive amount for it.

3. For each $j \in [m]$, let $s_j \in [n]$ be the smallest index i such that $\alpha'_i v_{ij} = h_j(\alpha')$, that is, s_j is the smallest index among the buyers who have the highest bid in good j .

4. For each $j \in [m]$, let $t_j \in [n]$ be the smallest index $i \neq s_j$ such that $\alpha'_i v_{ij} = \max_{k \neq s_j} \alpha'_k v_{kj}$ (so we have that $\alpha'_{t_j} v_{t_j j} = p_j(\alpha')$).

On the other hand, given any $I \subseteq [n]$, $W = (W_j \subseteq [n] : j \in [m])$, $s = (s_j \in [n] : j \in [m])$, and $t = (t_j \in [n] : j \in [m])$, we use $\text{LP}(I, W, s, t)$ to denote the following linear program on $n + nm + 1$ variables $\alpha = (\alpha_i : i \in [n])$, $q = (q_{ij} : i \in [n], j \in [m])$ and τ (where each variable q_{ij} captures the amount buyer i pays for good j):

$$\begin{aligned} & \text{minimize } \tau \\ & \tau \geq 0, \alpha_i \in [0, 1], q_{ij} \geq 0 \text{ for all } i \in [n] \text{ and } j \in [m] \\ & q_{ij} = 0 \text{ for all } j \in [m] \text{ and } i \notin W_j \\ & \alpha_{s_j} v_{s_j j} \geq \alpha_k v_{kj} \text{ for all } j \in [m] \text{ and } k \in [n] \\ & \alpha_{t_j} v_{t_j j} \geq \alpha_k v_{kj} \text{ for all } j \in [m] \text{ and } k \neq s_j \in [n] \\ (a) \quad & \alpha_i v_{ij} \geq \alpha_{s_j} v_{s_j j} \text{ for all } j \in [m] \text{ and } i \in W_j \\ (b) \quad & \sum_{k \in [n]} q_{kj} = \alpha_{t_j} v_{t_j j} \text{ for all } j \in [m] \\ (c) \quad & \sum_{j \in [m]} q_{ij} \leq B_i \text{ for all } i \in [n] \\ (d) \quad & \alpha_i \geq 1 - \tau \text{ for all } i \in I \text{ and } \sum_{j \in [m]} q_{ij} \geq (1 - \tau) B_i \text{ for all } i \notin I \end{aligned}$$

Here, (a) ensures that the buyers in W_j have the highest bid on good j ; (b) ensures that the total payment of all buyers for good j is equal to the second highest bid; (c) ensures that the budgets are satisfied; and (d) ensures that the not-too-much-unnecessary-pacing condition is satisfied. The lemma below follows directly from the definition of γ -approximate PE and the way I' , W' , s' and t' are extracted from (α', x') .

Lemma 11. Suppose (α', x') is a γ -approximate PE of G . Then (α', q', γ) is a feasible solution to the linear program $\text{LP}(I', W', s', t')$, where $q' = (q'_{ij})$ with $q'_{ij} = x'_{ij}p_j(\alpha')$.

On the other hand, the next lemma shows that if $\text{LP}(I, W, s, t)$ has a feasible solution $(\alpha, q, 0)$ for some I, W, s and t , then (α, x) is an exact pacing equilibrium, where $x = (x_{ij})$ and $x_{ij} = q_{ij}/p_j(\alpha)$ if $p_j(\alpha) > 0$; when $p_j(\alpha) = 0$ we set $x_{s,j} = 1$ and $x_{ij} = 0$ for all other i .

Lemma 12. If $(\alpha, q, 0)$ is a feasible solution to $\text{LP}(I, W, s, t)$, then (α, x) is an exact equilibrium.

Let $(\alpha, q, 0)$ be a feasible solution to $\text{LP}(I, W, s, t)$. Set α to be the pacing multipliers of buyers in G and define the allocation $x = (x_{ij})$ as above. Then, the LP constraints imply that the highest bid on good j is $h_j(\alpha) = \alpha_{s,j}v_{s,j}$ and the second highest bid is $p_j(\alpha) = \alpha_{t,j}v_{t,j}$. Next we note that, in the latter case, the constraints of the LP force the set of winners $\{i | x_{ij} > 0\}$ of good $j \in [m]$ to be a subset of W_j . This is because $x_{ij} > 0$ implies $q_{ij} > 0$ and $q_{ij} = 0$ for all $i \notin W_j$. Now, it is straightforward to see that constraints (a)-(d), in combination with $\tau = 0$, imply that (α, x) satisfies the corresponding conditions (a)-(d) of Definition 1. \square

Given the definition of $\text{LP}(I, W, s, t)$, there is a polynomial $r'(\cdot)$ such that

$$\max_{I, W, s, t} |\text{LP}(I, W, s, t)| \leq r'(|G|).$$

Now we can set γ to be smaller than $1/2^{r'(|G|)}$ (with bit length still polynomial in $|G|$). To finish the proof of Theorem 1, we let (α', x') be a γ -approximate PE of G . It follows from Lemma 11 that (α', q', γ) is a feasible solution to $\text{LP}(I', W', s', t')$. Next it follows from Fact 1 that this linear program has a feasible solution $(\alpha, q, 0)$ and the latter can be computed in polynomial time. Lemma 12 shows that (α, x) , which can be computed in polynomial time, is a pacing equilibrium of G .

5. Conclusion

We studied the computational complexity of pacing equilibria in second-price pacing games with multiplicative pacing. Our results show that finding a pacing equilibrium, whether exact or approximate, is a PPAD-complete problem. As discussed previously, these results close the open problem from Conitzer et al. [18] on the complexity of pacing equilibria, and make progress toward resolving the conjecture of Borgs et al. [9] by showing that their dynamics is unlikely to converge efficiently in second-price auctions. More generally, our results show that algorithms for budget-smoothing in auctions, an important problem for internet advertising, cannot be expected to efficiently find even approximate pacing equilibria in the worst case.

There are several interesting future questions and implications to investigate based on our work. Perhaps most importantly, we would like to understand exactly when budget-smoothing becomes hard. As discussed in the literature review (Balseiro and Gur [4]), developed regret minimization algorithms for the case of i.i.d. and continuous stochastic valuations. Yet our results imply that for general correlated valuations convergence cannot occur efficiently. The question is now which types of correlated stochastic valuations admit efficient algorithms, and which types are hard. It would also be interesting to understand whether other methods of budget smoothing (such as those discussed by Balseiro et al. [3]) lead to PPAD-complete equilibrium problems as well.

In the direction of positive results, our PPAD membership proof suggests that complementary pivoting may be a fruitful research direction for computing pacing equilibria. This is especially pertinent because approaches based on mixed-integer programming seem to scale poorly (Conitzer et al. [18]).

Acknowledgments

The authors thank the review team of Mathematics of Operations Research and the 22nd ACM Conference on Economics and Computation (EC'21) for their immensely valuable feedback that resulted in a significantly improved manuscript.

Appendix A. Proof of Theorem 3

Consider a $\{0, 1\}$ -cost $n \times n$ bimatrix game (A, B) and let $\epsilon = 1/n$. Recall that an ϵ -well-supported Nash equilibrium is a pair $(x, y) \in \Delta_n \times \Delta_n$ such that $x_i > 0$ for any $i \in [n]$ implies that $\sum_j A_{ij}y_j \leq \sum_j A_{kj}y_j + \epsilon$ for all k and $y_j > 0$ for any $j \in [n]$ implies $\sum_i x_i B_{ij} \leq \sum_i x_i B_{ik} + \epsilon$ for all $k \in [n]$.

In this section we show how to construct an SPP game G with $4n + 1$ buyers from the bimatrix game (A, B) in time polynomial in n such that every (δ, γ) -approximate PE of G , where $\delta = \gamma = \epsilon/n^6$, can be mapped back to an ϵ -well-supported Nash equilibrium of (A, B) in polynomial time. Theorem 2 follows from the PPAD-completeness of the problem of finding an ϵ -well-supported Nash equilibrium in a $\{0, 1\}$ -cost bimatrix game with $\epsilon = 1/n$ (Chen et al. [16]).

The SPP game G contains the following goods:

- Normalization goods: n goods $\{N(p, s)_1, \dots, N(p, s)_n\}$ for each $p \in \{1, 2\}$ and $s \in [n]$.
- Expenditure goods: n goods $\{E(p, s)_1, \dots, E(p, s)_n\}$ for each $p \in \{1, 2\}$ and $s \in [n]$.

- Threshold goods $T(p, s)$ for each $p \in \{1, 2\}$ and $s \in [n]$.

Set $v = 1/(16n)$. The set of buyers in G is defined as follows:

- Buyer $\mathbb{C}(p, s)$, $p \in \{1, 2\}$ and $s \in [n]$: $\mathbb{C}(p, s)$ has positive values for the following goods:
 - Normalization goods: $V(\mathbb{C}(p, s), N(p, s)_i) = 16$ for all $i \in [n] \setminus \{s\}$;
 $V(\mathbb{C}(p, s), N(p, s)_s) = 1$; and $V(\mathbb{C}(p, s), N(p, t)_s) = 1$ for all $t \in [n] \setminus \{s\}$.
 - Threshold good $T(p, s)$: $V(\mathbb{C}(p, s), T(p, s)) = 2n^4$.
 - Expenditure goods: $V(\mathbb{C}(p, s), E(p, s)_i) = 1$ for all $i \in [n]$.

For $p = 1$: $V(\mathbb{C}(1, s), E(2, t)_s) = vB_{st}$ for all $t \in [n]$.

For $p = 2$: $V(\mathbb{C}(2, s), E(1, t)_s) = vA_{st}$ for all $t \in [n]$.

For $p = 1$, the budget of buyer $\mathbb{C}(1, s)$ is $n/2 + n^4 + 1/4 + \sum_{t \in [n]} vA_{st}/2$;

For $p = 2$, the budget of buyer $\mathbb{C}(2, s)$ is $n/2 + n^4 + 1/4 + \sum_{t \in [n]} vB_{ts}/2$.

- Threshold Buyer \mathbb{T} : \mathbb{T} has positive values only for the following goods:

— Threshold goods: $V(\mathbb{T}, T(p, s)) = (1 - \delta)n^4$ for each $p \in \{1, 2\}$ and $s \in [n]$.

— Expenditure goods: $V(\mathbb{T}, E(1, s)_t) = vA_{st}/2$ and $V(\mathbb{T}, E(2, s)_t) = vB_{ts}/2$ for all $s, t \in [n]$.

Buyer \mathbb{T} has budget n^7 .

- Dummy buyer $\mathbb{D}(p, s)$, $p \in \{1, 2\}$ and $s \in [n]$: The budget of $\mathbb{D}(p, s)$ is v and she only values the normalization good $N(p, s)_s$: $V(\mathbb{D}(p, s), N(p, s)_s) = 1$.

Let \mathcal{E} be a (δ, γ) -approximate PE of the game G . We will use $\alpha(\cdot)$ to denote pacing multipliers of buyers in \mathcal{E} . Observe that, from the definition of approximate pacing equilibria, we must have $\alpha(\mathbb{T}) \in [1 - \gamma, 1]$. The following lemma establishes bounds on pacing multipliers of other buyers.

Lemma A.1. *For each $p \in \{1, 2\}$ and $s \in [n]$, we have*

$$\frac{(1 - \delta)^2}{2} \leq \alpha(\mathbb{C}(p, s)) \leq \frac{7}{8} \quad \text{and} \quad (1 - \delta) \cdot \alpha(\mathbb{C}(p, s)) \leq \alpha(\mathbb{D}(p, s)) \leq \frac{\alpha(\mathbb{C}(p, s))}{1 - \delta}.$$

Suppose for some $p \in \{1, 2\}$ and $s \in [n]$, we have $\alpha(\mathbb{C}(p, s)) < (1 - \delta)^2/2$. Then $\mathbb{C}(p, s)$ doesn't win any part of the threshold good $T(p, s)$. Observe that she has value at most 16 for every other good. Given that there are only $\tilde{O}(n^2)$ goods in G , she cannot possibly spend all her budget (which is $\Omega(n^4)$). This contradicts the assumption that \mathcal{E} is an approximate PE. Therefore, we have $\alpha(\mathbb{C}(p, s)) \geq (1 - \delta)^2/2$ for each $p \in \{1, 2\}$ and $s \in [n]$.

Next we prove the inequality about $\alpha(\mathbb{D}(p, s))$. Suppose $(1 - \delta)\alpha(\mathbb{D}(p, s)) > \alpha(\mathbb{C}(p, s))$ for some $p \in \{1, 2\}$ and $s \in [n]$. Then, $\mathbb{D}(p, s)$ wins all of good $N(p, s)_s$ at price $\alpha(\mathbb{C}(p, s)) \geq (1 - \delta)^2/2$. This violates her budget constraint and leads to a contradiction. Hence $\alpha(\mathbb{D}(p, s)) \leq \alpha(\mathbb{C}(p, s))/(1 - \delta)$. Moreover, if $\alpha(\mathbb{D}(p, s)) < (1 - \delta)\alpha(\mathbb{C}(p, s))$ (which implies $\alpha(\mathbb{D}(p, s)) < 1 - \delta = 1 - \gamma$) then her expenditure is zero. This violates the no unnecessary pacing condition. Hence the inequality about $\alpha(\mathbb{D}(p, s))$ must hold. Observe that, in particular, this means that the price of $N(p, s)_s$ is between $(1 - \delta)\alpha(\mathbb{C}(p, s))$ and $\alpha(\mathbb{C}(p, s))/(1 - \delta)$.

Finally suppose $\alpha(\mathbb{C}(p, s)) > 7/8$ for some $p \in \{1, 2\}$, $s \in [n]$. Then she wins:

- All of normalization good $N(p, s)_t$, for each $t \neq s$, by spending at least $(1 - \delta)^2/2$ on each of them because $\alpha(\mathbb{C}(p, t)) \geq (1 - \delta)^2/2$ by the first part of the proof.

- Part of normalization good $N(p, s)_s$ by spending at least $(1 - \delta)(7/8) - v$. This is because $N(p, s)_s$ has price at least $(1 - \delta)(7/8)$ and buyer $\mathbb{D}(p, s)$ only has budget v .

• All of threshold good $T(p, s)$ by spending $\alpha(\mathbb{T})(1 - \delta)n^4 \geq (1 - \delta)^2n^4$ (using $\gamma = \delta$).

• All of expenditure good $E(p, s)_t$, for each $t \in [n]$, by spending at least $\alpha(\mathbb{T})vA_{st}/2$ if $p = 1$ and $\alpha(\mathbb{T})vB_{ts}/2$ if $p = 2$.

Hence, the total expenditure of $\mathbb{C}(p, s)$ when $p = 1$ is at least

$$(1 - \delta)^2 \cdot \frac{n - 1}{2} + (1 - \delta) \cdot \frac{7}{8} - v + (1 - \delta)^2 n^4 + (1 - \delta) \sum_{t \in [n]} vA_{st}/2$$

which is strictly higher the budget (using $\delta = 1/n^7$). The same also holds for $p = 2$. In both cases, the budget constraint is violated, leading to a contradiction. Therefore, the lemma holds. \square

In particular, the above lemma implies that the total expenditure of each buyer $\mathbb{C}(p, s)$ is at least $(1 - \gamma)$ -fraction of her budget (and of course is also bounded from above by her budget). We also get the following corollary:

Corollary A.1. *For each $p \in \{1, 2\}$ and $s \in [n]$, the expenditure of $\mathbb{C}(p, s)$ on $N(p, s)_s$ lies in the following interval $[(1 - \delta)\alpha(\mathbb{C}(p, s)) - v, \alpha(\mathbb{C}(p, s)) - (1 - \delta)v]$*

Next, we define two vectors x' and y' with

$$x'_s = \{\alpha(1, s) - (\alpha(\mathbb{T})/2)\}^+ \quad \text{and} \quad y'_s = \{\alpha(2, s) - (\alpha(\mathbb{T})/2)\}^+$$

for each $s \in [n]$, where a^+ denotes $\max\{a, 0\}$. The following lemma will allow us to normalize x' and y' to obtain probability distributions.

Lemma A.2. *The following inequalities hold: $\sum_s x'_s > 1/8$ and $\sum_s y'_s > 1/8$.*

We prove $\sum_s x'_s > 1/8$. The proof of $\sum_s y'_s > 1/8$ is analogous. Suppose that $\sum_s x'_s \leq 1/8$. Then, buyer $\mathbb{B}(1,1)$ only wins a non-zero fraction of the following goods, and spends:

- At most $\alpha(\mathbb{C}(1,t))$ on each normalization good $N(1,t)_t$ for each $t \in [n]$. The total expenditure is

$$\sum_{t \in [n]} \alpha(\mathbb{C}(1,t)) \leq n\alpha(\mathbb{T})/2 + \sum_{t \in [n]} x'_t \leq n/2 + 1/8.$$

- At most $(1 - \delta)n^4$ on the threshold good $T(1,1)$.
- At most νA_{1t} on each expenditure good $E(1,1)_t$, $t \in [n]$.

Hence, the total expenditure of buyer $\mathbb{C}(1,1)$ is at most $n/2 + 1/8 + (1 - \delta)n^4 + \sum_t \nu A_{1t}$, which is strictly less than her budget, a contradiction. \square

Now, we are ready to define the mixed strategies (x, y) for the bimatrix game (A, B) . Set player 1's mixed strategy x to be $x_s = x'_s / \sum_i x'_i$ and player 2's mixed strategy y to be $y_s = y'_s / \sum_i y'_i$. These are valid mixed strategies because of Lemma 1 and Lemma 2. The next lemma shows that (x, y) is indeed an ϵ -well-supported Nash equilibrium of (A, B) .

Lemma A.3. *(x, y) is an ϵ -well-supported Nash equilibrium of the bimatrix game (A, B) .*

Assume there are $s, s^* \in [n]$ such that $x_s > 0$ but $\sum_t A_{st} y_t > \sum_t A_{s^*t} y_t + \epsilon$; the proof for y is analogous. Using $x_s > 0$, buyer $\mathbb{C}(1,s)$ spends non-zero amounts on the following goods:

- $\alpha(\mathbb{C}(1,t))$ on the normalization good $N(1,s)_t$ for each $t \neq s$.
- at least $(1 - \delta) \cdot \alpha(\mathbb{C}(1,s)) - \nu$ on the normalization good $N(1,s)_s$.
- $\alpha(\mathbb{T}) \cdot (1 - \delta)n^4$ on the threshold good $T(1,s)$.
- $\max\{\alpha(\mathbb{C}(2,t)), \alpha(\mathbb{T})/2\} \cdot \nu A_{st}$ on the expenditure good $E(1,s)_t$ for each $t \in [n]$.

Therefore, the total expenditure of buyer $\mathbb{C}(1,s)$ is at least

$$\begin{aligned} & \sum_{t \in [n]} \alpha(\mathbb{C}(1,t)) - \delta \cdot \alpha(\mathbb{C}(1,s)) - \nu + \alpha(\mathbb{T}) \cdot (1 - \delta)n^4 + \sum_{t \in [n]} \max\{\alpha(\mathbb{C}(2,t)), \alpha(\mathbb{T})/2\} \cdot \nu A_{st} \\ &= \sum_{t \in [n]} \alpha(\mathbb{C}(1,t)) - \delta \cdot \alpha(\mathbb{C}(1,s)) - \nu + \alpha(\mathbb{T}) \cdot (1 - \delta)n^4 + \alpha(\mathbb{T}) \sum_{t \in [n]} \nu A_{st}/2 + \nu \sum_{t \in [n]} y_t A_{st}. \end{aligned}$$

On the other hand, buyer $\mathbb{C}(1,s^*)$ spends (without assuming $x_{s^*} > 0$):

- $\alpha(\mathbb{C}(1,t))$ on the normalization good $N(1,s^*)_t$ for each $t \neq s^*$.
- at most $\alpha(\mathbb{C}(1,s^*)) - (1 - \delta)\nu$ on the normalization good $N(1,s^*)_{s^*}$.
- at most $\alpha(\mathbb{T}) \cdot (1 - \delta)n^4$ on the threshold good $T(1,s^*)$.
- $\max\{\alpha(\mathbb{C}(2,t)), \alpha(\mathbb{T})/2\} \cdot \nu A_{s^*t}$ on the expenditure good $E(1,s^*)_t$ for each $t \in [n]$.

Therefore, the total expenditure of buyer $\mathbb{C}(1,s^*)$ is at most

$$\sum_{t \in [n]} \alpha(\mathbb{C}(1,t)) - (1 - \delta)\nu + \alpha(\mathbb{T}) \cdot (1 - \delta)n^4 + \alpha(\mathbb{T}) \sum_{t \in [n]} \nu A_{s^*t}/2 + \nu \sum_{t \in [n]} y_t A_{s^*t}.$$

Using the assumption that $\sum_t A_{st} y_t > \sum_t A_{s^*t} y_t + \epsilon$, we have that the total expenditure of $\mathbb{C}(1,s)$ minus that of $\mathbb{C}(1,s^*)$, denoted by $(\ddagger 1)$, is at least

$$-\delta \cdot \alpha(\mathbb{C}(1,s)) - \delta\nu + \alpha(\mathbb{T}) \cdot \left(\sum_{t \in [n]} \nu A_{st}/2 - \sum_{t \in [n]} \nu A_{s^*t}/2 \right) + \epsilon\nu \geq \alpha(\mathbb{T}) \cdot \left(\sum_{t \in [n]} \nu A_{st}/2 - \sum_{t \in [n]} \nu A_{s^*t}/2 \right) + \epsilon\nu/2$$

using $\epsilon\nu \gg \delta$. On the other hand, the budget of $\mathbb{C}(1,s)$ minus that of $\mathbb{C}(1,s^*)$, denoted $(\ddagger 2)$, is

$$\sum_{t \in [n]} \nu A_{st}/2 - \sum_{t \in [n]} \nu A_{s^*t}/2.$$

Using $\alpha(\mathbb{T}) \geq 1 - \gamma$ and $\gamma = 1/n^7$, we have $(\ddagger 1) \geq (\ddagger 2) + \epsilon\nu/3$. However, the total expenditure of $\mathbb{C}(1,s)$ is at most her budget and the total expenditure of $\mathbb{C}(1,s^*)$ is at least $(1 - \gamma)$ -fraction of her budget. Given that the budget of $\mathbb{C}(1,s^*)$ is $O(n^4)$, we also have

$$(\ddagger 1) \leq (\ddagger 2) + \gamma \cdot O(n^4) = (\ddagger 2) + O(1/n^3),$$

a contradiction because $\epsilon\nu = \Omega(1/n^2)$. \square

Theorem 3 follows from the PPAD-hardness of finding an ϵ -well-supported Nash equilibrium in a $\{0, 1\}$ -cost bimatrix game (Chen et al. [16]).

Appendix B. Proof of Claim 1

Before stating the proof of Claim 1, we state and prove the following useful lemma.

Lemma B.1. *If $\beta \in S$ is labelled i , then $\beta_i \geq \min\left\{\frac{1}{n}, \frac{B_{\min}}{2nv_{\max}}\right\}$.*

Without loss of generality, we will prove the lemma for $i = 1$. Suppose $\beta \in S$ is labelled 1 according to the above procedure. First, $\beta_1 > 0$ follows as a direct consequence. Furthermore, as $\max_i \beta_i \geq 1/n$ and $\sum_i \beta_i = 1$, we get $t^*(\beta) = t_1 \leq n$. We consider the two possible binding cases which can define t_1 . If $t_1 = 1/\beta_1$, then $\beta_1 \geq 1/n$, and thus the lemma holds. On the other hand, if $t_1 = \frac{B_1}{\sum_j x_{1j} p_j(\beta)}$, then $\sum_j x_{1j} p_j(\beta) > 0$ and

$$B_1 = t_1 \sum_j x_{1j} p_j(\beta) \leq n \sum_{j: x_{1j} > 0} \max_i \beta_i v_{1j} \leq n \sum_{j: x_{1j} > 0} \frac{\beta_1 v_{1j}}{(1 - \delta)} \leq n \sum_j \frac{\beta_1 v_{1j}}{(1 - \delta)}$$

where the second inequality follows from the definition of (δ, γ) -approximate pacing equilibrium. Therefore, $\beta_1 \geq \min\left\{\frac{1}{n}, \frac{(1 - \delta)B_1}{\sum_j v_{1j}}\right\}$. \square

Proof of Claim 1. Let $B_{\min} = \min_{i \in [n]} B_i$, $B_{\max} = \max_{i \in [n]} B_i$, $v_{\max} = \max_{i,j} v_{ij}$ and $v_{\min} = \min_{i,j: v_{ij} > 0} v_{ij}$. In this proof, we will use the following facts: if f, g are Lipschitz functions with Lipschitz constants L_f, L_g , then

- (a) $f + g$ is Lipschitz with constant $L_f + L_g$
- (b) $\max\{f, g\}$ is Lipschitz with constant $\max\{L_f, L_g\}$.
- (c) If $|f|, |g| \leq M$, then fg is Lipschitz with constant $M(L_f + L_g)$.

Define $y_{ij} : S \rightarrow \mathbb{R}$ as $y_{ij}(\beta) = [\beta_i v_{ij} - (1 - \delta) \max_k \beta_k v_{kj}]^+$. Using facts (a) and (b), we can write

$$|y_{ij}(\beta) - y_{ij}(\beta')| \leq 2v_{\max} \|\beta - \beta'\|_{\infty}$$

Consider $\beta \in S_0$ and $i \in [n]$. As S_0 is panchromatic, there exists $\beta' \in S_0$ such that $T(\beta') = i$. By Lemma B.1, we get

$$\beta'_i \geq \min\left\{\frac{1}{n}, \frac{B_{\min}}{2nv_{\max}}\right\}$$

Then, using the definition of ω , we get the following equivalent statements:

$$\beta_i \geq \frac{1}{2} \min\left\{\frac{1}{n}, \frac{B_{\min}}{2nv_{\max}}\right\} \iff \frac{1}{\beta_i} \leq U := 2 \max\left\{n, \frac{2nv_{\max}}{B_{\min}}\right\}$$

Hence, for $\beta, \beta' \in S_0$, we have

$$\begin{aligned} \left| \frac{1}{\sum_r y_{rj}(\beta)} - \frac{1}{\sum_r y_{rj}(\beta')} \right| &= \left| \frac{\sum_r y_{rj}(\beta') - \sum_r y_{rj}(\beta)}{\sum_r y_{rj}(\beta) \sum_r y_{rj}(\beta')} \right| \\ &\leq \frac{2nv_{\max} U^2}{\delta^2 v_{\min}^2} \cdot \|\beta - \beta'\|_{\infty} \end{aligned}$$

Using fact (c), for $\beta, \beta' \in S_0$, we can write

$$|x_{ij}(\beta) - x_{ij}(\beta')| \leq \max\left\{v_{\max}, \frac{U}{\delta v_{\min}}\right\} \left[2v_{\max} + \frac{2nv_{\max} U^2}{\delta^2 v_{\min}^2} \right] \cdot \|\beta - \beta'\|_{\infty}$$

Set $\bar{U} = \max\left\{v_{\max}, \frac{U}{\delta v_{\min}}\right\} \left[2v_{\max} + \frac{2nv_{\max} U^2}{\delta^2 v_{\min}^2} \right]$. Also, note that for $\beta, \beta' \in S$,

$$|p_j(\beta) - p_j(\beta')| \leq v_{\max} \|\beta - \beta'\|_{\infty}$$

For $\beta, \beta' \in S_0$, combining the above Lipschitz conditions using facts (a) and (c) yields

$$\left| \sum_j x_{ij}(\beta) p_j(\beta) - \sum_j x_{ij}(\beta') p_j(\beta') \right| \leq mv_{\max} (\bar{U} + v_{\max}) \|\beta - \beta'\|_{\infty}$$

Set $W := mv_{\max} (\bar{U} + v_{\max})$. Define

$$P^* := \left\{ i \in [n] \mid \exists \beta \in T \text{ s.t. } \frac{B_i}{\sum_j x_{ij}(\beta) p_j(\beta)} < \frac{1}{\beta_i} \right\}$$

For $i \in P^*$ and $\beta \in S_0$, we can write $\frac{B_i}{\sum_j x_{ij}(\beta) p_j(\beta)} < \frac{1}{\beta_i} \leq U$, which implies $\frac{1}{\sum_j x_{ij}(\beta) p_j(\beta)} \leq \frac{U}{B_{\min}}$.

Therefore, for $\beta, \beta' \in S_0$ and $i \in P^*$, we have

$$\left| \frac{B_i}{\sum_j x_{ij}(\beta) p_j(\beta)} - \frac{B_i}{\sum_j x_{ij}(\beta') p_j(\beta')} \right| \leq B_{\max} \frac{U^2}{B_{\min}^2} W \|\beta - \beta'\|_{\infty} \leq \frac{B_{\max} U^2 W}{B_{\min}^2} \|\beta - \beta'\|_{\infty}$$

Also, for $\beta, \beta' \in S_0$ and $i \in [n]$, we have

$$\left| \frac{1}{\beta_i} - \frac{1}{\beta'_i} \right| \leq U^2 \|\beta - \beta'\|_\infty$$

Note that for $\beta \in T$, we can rewrite $t^*(\beta)$ as follows

$$t^*(\beta) = \min \left\{ \min_{i \in [n]} \frac{1}{\beta_i}, \min_{i \in P^*} \min \left\{ \frac{1}{\beta_i}, \frac{B_i}{\sum_j x_{ij}(\beta) p_j(\beta)} \right\} \right\}$$

Using fact (b), for $\beta, \beta' \in T$,

$$|t^*(\beta) - t^*(\beta')| \leq 2n \max \left\{ U^2, \frac{B_{\max} U^2 L}{B_{\min}^2} \right\} \|\beta - \beta'\|_\infty$$

Therefore, for $i \in [n]$, total payment made by buyer i is Lipschitz for $\beta \in S_0$:

$$\begin{aligned} & \left| \sum_j x_{ij}(\beta) t^*(\beta) p_j(\beta) - \sum_j x_{ij}(\beta') t^*(\beta') p_j(\beta') \right| \\ & \leq \max\{nv_{\max}, n\} \left(W + 2n \max \left\{ U^2, \frac{B_{\max} U^2 W}{B_{\min}^2} \right\} \right) \|\beta - \beta'\|_\infty \end{aligned}$$

Hence, the claim holds because

$$\max\{nv_{\max}, n\} \left(W + 2n \max \left\{ U^2, \frac{B_{\max} U^2 W}{B_{\min}^2} \right\} \right) \leq L = \left(\frac{2^{|G|}}{\delta} \right)^{10,000}$$

Appendix C. Incorporating Reserve Prices

Consider the setting in which each item j has a reserve price r_j . Now, a buyer wins a good j only if her bid is the highest bid $h_j(\alpha)$ and it is greater than or equal to the reserve r_j . Moreover, the price of good j is the maximum of the second highest bid $p_j(\alpha)$ and its reserve price r_j . In the presence of reserve prices, we will use $H_j(\alpha) := \max\{h_j(\alpha), r_j\}$ to denote the winning threshold of good j and $P_j(\alpha) := \max\{p_j(\alpha), r_j\}$ to denote the price of good j . The next example illustrates that one needs to be careful in the way one extends the definition of pacing equilibrium (Definition 1) to model the presence of reserves.

Example C.1. There is one buyer and one good. The buyer values the good at 4 and has a budget of 1. The goods has a reserve price of 2. If she bids strictly less than 1/2, then she does not win any part of the good. On the other hand, if we assume that she wins the entire good upon bidding 1/2 or higher, then she violates her budget upon doing so. This suggests that a pacing equilibrium might not even exist if we extend it naively to the setting with reserves. Instead, we will take the approach that, in a pacing equilibrium, the seller may decide to not sell a fraction of a good if the highest bid is equal to the reserve price of that good. With this new definition, we can see that a pacing equilibrium does in fact exist, namely, when the buyer has a pacing multiplier of 1/2 and wins 1/2 of the item.

Inspired by the above example, we define pacing equilibrium for the setting with reserves.

Definition C.1. (Pacing Equilibria with Reserves). Given an SPP game with reserves $G = (n, m, (v_{ij}), (B_i), (r_j))$, we say (α, x) with $\alpha = (\alpha_i) \in [0, 1]^n$, $x = (x_{ij}) \in [0, 1]^{nm}$ and $\sum_{i \in [n]} x_{ij} \leq 1$ for all $j \in [m]$ is a *pacing equilibrium* if

- (a) Only buyers above the winning threshold win the good: $x_{ij} > 0$ implies $\alpha_i v_{ij} = H_j(\alpha)$.
- (b) Full allocation of each good for which the highest bid exceeds the reserve price: $h_j(\alpha) > r_j$ implies $\sum_{i \in [n]} x_{ij} = 1$.
- (c) Budgets are satisfied: $\sum_{j \in [m]} x_{ij} P_j(\alpha) \leq B_i$.
- (d) No unnecessary pacing: $\sum_{j \in [m]} x_{ij} P_j(\alpha) < B_i$ implies $\alpha_i = 1$.

Next, we extend our PPAD-membership result to the setting with reserves.

Theorem C.1. *Finding a pacing equilibrium in a SPP game with reserves is in PPAD.*

Consider a pacing game with reserve prices G and the corresponding pacing game without reserve prices G' . Add an auxiliary buyer a to G' who values good j at r_j for all $j \in [m]$ and has a budget large enough to ensure that her pacing multiplier is always 1 in every pacing equilibrium (this can be achieved by setting her budget to be the sum of all values $\{v_{ij}\}$ and reserve prices $\{r_j\}$). We will call this updated game G'_+ . The theorem follows from the simple observation that if we find a pacing equilibrium (α, x) for G'_+ and disregard the terms corresponding to the auxiliary buyer, then we get a pacing equilibrium (α_{-a}, x_{-a}) for G . This is because, in any pacing equilibrium of G'_+ , the auxiliary buyer has a multiplier of 1 and hence bids r_j on good j for all $j \in [m]$. Moreover, any amount that the auxiliary buyer wins in (α, x) can be thought of as being not sold by the seller. As (α, x) satisfies Definition 1, it is straightforward to check that (α_{-a}, x_{-a}) satisfies Definition C.1. \square

We conclude this section by noting that our hardness results extend directly to the setting with reserves because it reduces to the setting without reserves when $r_j = 0$ for all goods $j \in [m]$.

Appendix D. Perturbed Second-Price Pacing Games

Before stating and proving the results, we define the relevant equilibrium notions. For a perturbed pacing game $(n, m, (v_{ij}), (B_i), \delta)$, let $p'_{ij}(\alpha)$ denote the expected payment made by buyer i on good j when the buyers use multipliers $\alpha \in [0, 1]^n$. Moreover, let $x_{ij}(\alpha)$ be the probability of buyer i winning good j when the buyers use the multipliers α .

Definition D.1. Consider a perturbed SPP game $(n, m, (v_{ij}), (B_i), \delta)$. Then, $\alpha \in [0, 1]^n$ is a pacing equilibrium of the perturbed SPP if:

- Budgets are satisfied: $\sum_{j=1}^m p'_{ij}(\alpha) \leq B_i$
- No unnecessary pacing: If $\sum_{j=1}^m p'_{ij}(\alpha) < B_i$, then $\alpha_i = 1$

Moreover, $\alpha \in [0, 1]^n$ is a γ -approximate pacing equilibrium of the perturbed SPP if:

- Budgets are satisfied: $\sum_{j=1}^m p'_{ij}(\alpha) \leq B_i$
- Not too much unnecessary pacing: If $\sum_{j=1}^m p'_{ij}(\alpha) < (1 - \gamma)B_i$, then $\alpha_i \geq (1 - \gamma)v_{ij}$

Theorem D.1. Computing a γ -approximate pacing equilibrium of a perturbed SPP game $(n, m, (v_{ij}), (B_i), \delta)$ is PPAD-hard when $\delta = \gamma = 1/n^8$.

First observe that

$$(1 - \gamma)(1 - \delta) = (1 - n^{-8})^2 = 1 + n^{-16} - 2n^{-8} \geq 1 - n^{-7}$$

We will prove the theorem by reducing from the problem of computing approximate pacing equilibria of SPP games. Consider an SPP game $G = (n, m, (v_{ij}), (B_i))$. Define a perturbed SPP game $G' = (n, m, (v_{ij}), (B'_i), \delta)$ such that $B'_i = (1 - \delta)B_i$. Let α be a γ -approximate pacing equilibrium of the perturbed SPP game G' . Then, as $\epsilon_{ij} \in [1 - \delta, 1]$, we get that

$$(1 - \delta)x_{ij}(\alpha)p_j(\alpha) \leq p'_{ij}(\alpha) \leq x_{ij}(\alpha)p_j(\alpha) \quad \forall i \in [n], j \in [m] \quad (\text{D.1})$$

where, as earlier, $p_j(\alpha)$ denotes the second highest bid in an SPP game when the buyers use multipliers α . To complete the proof, it suffices to show that $(\alpha, x(\alpha))$ is a (δ, γ') -approximate pacing equilibrium of the SPP game G for $\gamma' = 1/n^7$. We establish the required properties below:

- (a) As $\epsilon_{ij} \in [1 - \delta, 1]$, $x_{ij}(\alpha) > 0$ only if $\alpha_i v_{ij} \geq (1 - \delta)\max_{k \in [n]} \alpha_k v_{kj}$
- (b) Full allocation of each good with positive bid: This follows directly from the allocation rules of a second-price auction.
- (c) Budgets are satisfied: α being budget feasible for the perturbed SPP game G implies

$$\sum_{j=1}^m p'_{ij}(\alpha) \leq B'_i = (1 - \delta)B_i$$

for all $i \in [n]$. As $p'_{ij}(\alpha) \geq (1 - \delta)x_{ij}(\alpha)p_j(\alpha)$, we get $\sum_{j=1}^m x_{ij}(\alpha)p_j(\alpha) \leq B_i$ as required.

- (d) Not too much unnecessary pacing: Suppose $\sum_{j=1}^m x_{ij}(\alpha)p_j(\alpha) < (1 - \gamma')B_i$ for some buyer $i \in [n]$. Then, using (D.1), we get

$$\sum_{j=1}^m p'_{ij}(\alpha) < \frac{(1 - \gamma')}{(1 - \delta)} \cdot (1 - \delta)B_i = \frac{(1 - \gamma')}{(1 - \delta)} B'_i \leq (1 - \gamma)B'_i$$

where we have used $(1 - \gamma)(1 - \delta) \geq (1 - n^{-7}) = (1 - \gamma')$. Now, as α is a γ -approximate equilibrium of the perturbed SPP game G' , we get $\alpha_i \geq 1 - \gamma \geq 1 - n^{-7} = 1 - \gamma'$.

Hence, we have shown that $(\alpha, x(\alpha))$ is a (δ, γ') -approximate pacing equilibrium for the SPP game G , where $\delta \leq n^{-7}$ and $\gamma' = n^{-7}$. As the perturbed SPP game G' can be constructed from the SPP game G in polynomial time, the theorem follows from Theorem 3. \square

Let the expected utility of buyer i in a perturbed SPP game under multipliers α be denoted by $u_i(\alpha)$, i.e.,

$$u_i(\alpha) = \mathbb{E}_{\{\epsilon_{ij}\}_{i,j}} \left[\sum_{j=1}^m \left(v_{ij}\epsilon_{ij} - \max_{k \neq i} \alpha_k v_{kj}\epsilon_{kj} \right) \mathbf{1} \left(\alpha_i v_{ij}\epsilon_{ij} \geq \max_{k \neq i} \alpha_k v_{kj}\epsilon_{kj} \right) \right]$$

Definition D.2. Consider a perturbed SPP game $(n, m, (v_{ij}), (B_i), \delta)$. A vector of pacing multipliers α is called a Nash equilibrium of this game if for each $i \in [n]$ and α'_i such that $\sum_{j=1}^m p'_{ij}(\alpha'_i, \alpha_{-i}) \leq B_i$, we have $u_i(\alpha_i, \alpha_{-i}) \geq u_i(\alpha'_i, \alpha_{-i})$.

Lemma D.1. Consider a perturbed SPP game $(n, m, (v_{ij}), (B_i), \delta)$ and let α be a Nash equilibrium of this game. If $\sum_{j=1}^m p'_{ij}(\alpha) < B_i$ and $\alpha_i < 1$, then $\sum_{j=1}^m p'_{ij}(\alpha) = \sum_{j=1}^m p'_{ij}(1, \alpha_{-i})$.

Suppose α is a Nash equilibrium of the game but not a pacing equilibrium, and buyer i satisfies $\sum_{j=1}^m p'_{ij}(\alpha) < B_i$ and $\alpha_i < 1$. For contradiction, suppose $\sum_{j=1}^m p'_{ij}(\alpha) < \sum_{j=1}^m p'_{ij}(1, \alpha_{-i})$. Now, as the distribution of ϵ_{ij} is continuous, $x \mapsto p_{ij}(x, \alpha_{-i})$

is a continuous non-decreasing function. By the Intermediate Value Theorem, there exists $\alpha_i^* \in (\alpha_i, 1)$ such that

$$\sum_{j=1}^m p'_{ij}(\alpha_i^*, \alpha_{-i}) \leq B_i.$$

Now, observe that buyer i wins good j if and only if

$$\alpha_i^* v_{ij} \epsilon_{ij} \geq \max_{k \neq i} \alpha_k v_{kj} \epsilon_{kj}$$

Therefore, $v_{ij} \epsilon_{ij} \geq p'_{ij}(\alpha_i^*, \alpha_{-i}) / \alpha_i^*$. As $\alpha_i^* < 1$, we get that

$$u_i(\alpha_i^*, \alpha_{-i}) - u_i(\alpha_i, \alpha_{-i}) \geq \frac{1}{\alpha_i^*} \cdot \left[\sum_{j=1}^m p'_{ij}(\alpha_i^*, \alpha_{-i}) - \sum_{j=1}^m p'_{ij}(\alpha_i, \alpha_{-i}) \right] > 0$$

This contradicts the fact that α is a Nash equilibrium. Hence, the Lemma holds. \square

Corollary D.1. Consider a perturbed SPP game $(n, m, (v_{ij}), (B_i), \delta)$ and let α be a Nash equilibrium of this game. If $\sum_{j=1}^m p'_{ij}(1, \alpha_{-i}) > \sum_{j=1}^m p'_{ij}(\alpha)$, then we have $\sum_{j=1}^m p'_{ij}(\alpha) = B_i$. Furthermore, as a consequence, if $\sum_{j=1}^m p'_{ij}(1, \alpha_{-i}) > B_i$, then $\sum_{j=1}^m p'_{ij}(\alpha) = B_i$.

Theorem D.2. Computing a Nash equilibrium of a perturbed SPP game $(n, m, (v_{ij}), (B_i), \delta)$ is PPAD-hard when $\delta = 1/n^8$.

Let G be the SPP game constructed in Appendix A for the proof of Theorem 3. Like the proof of Theorem D.1, define a perturbed SPP game $G' = (n, m, (v_{ij}), (B'_i), \delta)$ such that $B'_i = (1 - \delta)B_i$. Moreover, define an auxiliary perturbed SPP game $G'' = (n + 1, m + 1, (v_{ij}), (B'_i), \delta)$ by adding one more buyer and one more good to G' . We denote the new buyer by \mathbb{T}^* and the new good by S . Buyer \mathbb{T}^* has value 1 for good S , that is, $V(\mathbb{T}^*, S) = 1$ and does not value any other good. She has a budget of n^7 (large enough to never be binding). The only other buyer who has a nonzero value for S is the Threshold buyer \mathbb{T} , who has a value of 1, that is, $V(\mathbb{T}, S) = 1$.

We begin by showing that every Nash equilibrium of G'' is also a pacing equilibrium. Let α be a Nash equilibrium of G'' . As a first step, we show that $\alpha(\mathbb{T}) = \alpha(\mathbb{T}^*) = 1$. We do so by ruling out the other cases:

1. If $\alpha(\mathbb{T}) < \alpha(\mathbb{T}^*)$, then buyer \mathbb{T} can strictly increase her utility by setting $\alpha(\mathbb{T}) = 1$ as this allows her to win a strictly larger fraction of good S .
2. Similarly, if $\alpha(\mathbb{T}^*) < \alpha(\mathbb{T})$, then buyer \mathbb{T}^* can strictly increase her utility by setting $\alpha(\mathbb{T}^*) = 1$ as this allows her to win a strictly larger fraction of good S .
3. If $\alpha(\mathbb{T}) = \alpha(\mathbb{T}^*) < 1$, then buyer \mathbb{T} can strictly increase her utility by setting $\alpha(\mathbb{T}) = 1$ as this allows her to win a strictly larger fraction of good S .

For every other buyer in G'' , we use Corollary D.1 to show that they exactly spend their budget.

If $\alpha(\mathbb{C}(p, s)) \leq (1 - \delta)/2$, then the buyer $\mathbb{C}(p, s)$ wins no part of the threshold good $T(p, s)$ and spends strictly less than her budget because she has value at most 16 for all of the other goods and there are at most $O(n^2)$ such goods compared with her budget which is $\Omega(n^2)$. On the other hand, she can win all of the threshold good $T(p, s)$ by setting $\alpha(\mathbb{C}(p, s)) = 1$ and spend strictly more. Hence, by Corollary D.1, we get that she exactly spends her budget, which is a contradiction. Therefore, $\alpha(\mathbb{C}(p, s)) \geq (1 - \delta)/2$.

Consider a dummy buyer $\mathbb{D}(p, s)$. If we set $\alpha(\mathbb{D}(p, s)) = 1$, then she wins at least half of the normalization good $N(p, s)_s$ at a price of at least $\alpha(\mathbb{C}(p, s))$ which violates her budget of $1/(16n)$. Thus, Corollary D.1 implies that she exactly spends her budget under the Nash equilibrium α .

Consider buyer $\mathbb{C}(p, s)$. If we set $\alpha(\mathbb{C}(p, s)) = 1$, she wins:

- All of normalization good $N(p, s)_t$, for each $t \neq s$, by spending at least $(1 - \delta)/2$ on each of them because $\alpha(\mathbb{C}(p, t)) \geq (1 - \delta)/2$ by the earlier part of the proof.
- Part of normalization good $N(p, s)_s$ by spending at least $(1 - \delta) - v$. This is because $N(p, s)_s$ has price at least $(1 - \delta)$ and buyer $\mathbb{D}(p, s)$ only has budget v .
- All of threshold good $T(p, s)$ by spending at least $\alpha(\mathbb{T})(1 - \delta)n^4 = (1 - \delta)n^4$.
- All of expenditure good $E(p, s)_t$, for each $t \in [n]$, by spending at least $\alpha(\mathbb{T})vA_{st}/2$ if $p = 1$ and $\alpha(\mathbb{T})vB_{ts}/2$ if $p = 2$.

Hence, the total expenditure of $\mathbb{C}(p, s)$ when $p = 1$ is at least

$$(1 - \delta) \cdot \frac{n - 1}{2} + (1 - \delta) \cdot -v + (1 - \delta)n^4 + \sum_{t \in [n]} vA_{st}/2$$

which is strictly higher than her budget. Similar statement holds for $p = 2$. Therefore, Corollary D.1 implies that buyer $\mathbb{C}(p, s)$ exactly spends her budget.

Hence, we have shown that every buyer either has her multiplier equal to 1 or exactly spends her budget, which means that α is a pacing equilibrium. Moreover, from our construction of G'' from G' , we get that the restriction of α to the buyers other than \mathbb{T}^* is a pacing equilibrium for the game G' . This is because only the Threshold buyer \mathbb{T} is affected by this change and her multipliers satisfies $\alpha(\mathbb{T}) = 1$ and she spends strictly less than her budget. Finally, as we showed

in the proof of Theorem D.1, $(\alpha, x(\alpha))$ is a (δ, γ) -approximate pacing equilibrium of the SPP game G where $\delta = \gamma = 1/n^7$. Invoking Theorem 3 completes the proof. \square

Endnotes

¹ See <https://www.facebook.com/business/help/1754368491258883?id=561906377587030>.

² Bidding your value is a dominant strategy in second price auctions without budgets.

³ While we cite the 2021 journal version of that paper, the conjecture was made first in the 2017 arXiv version of that paper. It was also published in their 2018 conference version of the paper, which appeared at the Conference on Web and Internet Economics (WINE) that year.

⁴ One can imagine a strategy where a buyer consistently bids higher in an attempt to run other buyers out of budget early with the goal of winning goods for cheap later on. Although interesting, an analysis of these strategies would require studying a complicated incomplete-information extensive-form game, which is not the focus of this work. Moreover, in practice, these pacing algorithms are predominantly implemented by platforms who have no incentive to take advantage of some advertisers on behalf of other ones. Finally, while such strategies may seem appealing in toy examples, in a large-scale market, where the budget of an individual advertiser is small relative to the whole market, such an approach is unlikely to be possible.

⁵ The circuit has $O(n \log(1/\omega))$ input variables to encode a point of S_ω and has $\lceil \log n \rceil$ output gates to encode the output of the Sperner coloring T .

References

- [1] Ashlagi I, Braverman M, Hassidim A, Lavi R, Tennenholtz M (2010) Position auctions with budgets: Existence and uniqueness. *The BE J. Theoret. Econom.* 10(1):1–32.
- [2] 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 in Theoret. Comput. Sci. Conf. (ITCS 2021)*, Leibniz International Proceedings in Informatics (LIPIcs), vol. 185 (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany), 84:1–84:1.
- [3] Balseiro S, Kim A, Mahdian M, Mirrokni V (2021) Budget-management strategies in repeated auctions. *Oper. Res.* 69(3):859–876.
- [4] Balseiro SR, Gur Y (2019) Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Sci.* 65(9):3952–3968.
- [5] Balseiro SR, Besbes O, Weintraub GY (2015) Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Sci.* 61(4):864–884.
- [6] Balseiro SR, Kroer C, Kumar R (2022) Contextual standard auctions with budgets: Revenue equivalence and efficiency guarantees. *Proc. 23rd ACM Conf. Econom. Comput.* (ACM, New York), 476–476.
- [7] Bei X, Garg J, Hoefer M, Mehlhorn K (2016) Computing equilibria in markets with budget-additive utilities. Sankowski P, Zaroliagis C, eds. *24th Annual Eur. Sympos. Algorithms (ESA 2016)*, Leibniz International Proceedings in Informatics (LIPIcs), vol. 57 (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany), 8:1–8:14.
- [8] Bitansky N, Paneth O, Rosen A (2015) On the cryptographic hardness of finding a Nash equilibrium. Guruswami V, ed. *IEEE 56th Annual Symp. Foundations Comput. Sci. (FOCS)* (IEEE, Piscataway, NJ), 1480–1498.
- [9] Borgs C, Chayes J, Immorlica N, Jain K, Etesami O, Mahdian M (2007) Dynamics of bid optimization in online advertisement auctions. *Proc. 16th Internat. Conf. World Wide Web* (ACM, New York), 531–540.
- [10] Budish E (2011) The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *J. Political Econom.* 119(6):1061–1103.
- [11] Chen X, Deng X (2006) Settling the complexity of two-player Nash equilibrium. *2006 47th Annual IEEE Symp. Foundations Comput. Sci. (FOCS'06)* (IEEE, Piscataway, NJ), 261–272.
- [12] Chen X, Deng X (2009) On the complexity of 2d discrete fixed point problem. *Theoret. Comput. Sci.* 410(44):4448–4456.
- [13] Chen X, Teng SH (2009) Spending is not easier than trading: On the computational equivalence of Fisher and Arrow-Debreu equilibria. Dong Y, Du D-Z, eds. *Algorithms and Computation* (Springer, Berlin, Heidelberg), 647–656.
- [14] Chen X, Kroer C, Kumar R (2021) Throttling equilibria in auction markets. Feldman M, Fu H, Talgam-Cohen I, eds. *Web and Internet Econom. – 17th Internat. Conf. (WINE)*, Lecture Notes in Computer Science, vol. 13112 (Springer, Cham, Switzerland), 551.
- [15] Chen X, Paparas D, Yannakakis M (2017) The complexity of non-monotone markets. *J. ACM.* 64(3):1–56.
- [16] Chen X, Teng SH, Valiant P (2007) The approximation complexity of win-lose games. *Proc. 18th Annual ACM-SIAM Symp. Discrete Algorithms* (SIAM, Philadelphia), 159–168.
- [17] Choudhuri AR, Hubáček P, Kamath C, Pietrzak K, Rosen A, Rothblum GN (2019) Finding a Nash equilibrium is no easier than breaking Fiat-Shamir. *Proc. 51st Annual ACM SIGACT Symp. Theory Comput.* (ACM, New York), 1103–1114.
- [18] Conitzer V, Kroer C, Sodomka E, Stier-Moses NE (2022a) Multiplicative pacing equilibria in auction markets. *Oper. Res.* 70(2):963–989.
- [19] Conitzer V, Kroer C, Panigrahi D, Schrijvers O, Stier-Moses NE, Sodomka E, Wilkens CA (2022b) Pacing equilibrium in first price auction markets. *Management Sci.* 68(12):8515–8535.
- [20] Daskalakis C, Goldberg PW, Papadimitriou CH (2009) The complexity of computing a Nash equilibrium. *SIAM J. Comput.* 39(1):195–259.
- [21] Deng X, Qi Q, Saberi A (2012) Algorithmic solutions for envy-free cake cutting. *Oper. Res.* 60(6):1461–1476.
- [22] Dobzinski S, Leme RP (2014) Efficiency guarantees in auctions with budgets. Esparza J, Fraigniaud P, Husfeldt T, Koutsoupias E, eds. *Automata, Languages, and Programming* (Springer, Berlin, Heidelberg), 392–404.
- [23] Dobzinski S, Lavi R, Nisan N (2012) Multi-unit auctions with budget limits. *Games Econom. Behav.* 74(2):486–503.
- [24] Etessami K, Yannakakis M (2010) On the complexity of Nash equilibria and other fixed points. *SIAM J. Comput.* 39(6):2531–2597.
- [25] Facebook (2017) Your guide to Facebook bid strategy. Accessed June 30, 2021, https://www.facebook.com/gms_hub/share/biddingstrategyguide_final.pdf.
- [26] Fearnley J, Goldberg PW, Hollender A, Savani R (2021) The complexity of gradient descent: $\text{CLS} = \text{PPAD} \cap \text{PLS}$. *J. ACM* 70(1):1–74.

- [27] Filos-Ratsikas A, Hollender A, Sotiraki K, Zampetakis M (2023) Consensus-halving: Does it ever get easier? *SIAM J. Comput.* 52(2):412–451.
- [28] Filos-Ratsikas A, Giannakopoulos Y, Hollender A, Lazos P, Poças D (2023) On the complexity of equilibrium computation in first-price auctions. *SIAM J. Comput.* 52(1):80–131.
- [29] Garg S, Pandey O, Srinivasan A (2016) Revisiting the cryptographic hardness of finding a Nash equilibrium. *Proc. 36th Annual Internat. Cryptology Conf. Adv. Cryptol.* (Springer, Berlin, Heidelberg), 579–604.
- [30] Goel G, Mirrokni V, Leme RP (2015) Polyhedral clinching auctions and the adwords polytope. *J. ACM* 62(3):1–27.
- [31] Goldberg PW (2011) A survey of PPAD-completeness for computing Nash equilibria. Preprint, submitted March 14, <https://arxiv.org/abs/1103.2709>.
- [32] Hubáček P, Yodov E (2020) Hardness of continuous local search: Query complexity and cryptographic lower bounds. *SIAM J. Comput.* 49(6):1128–1172.
- [33] Kuhn HW (1960) Some combinatorial lemmas in topology. *IBM J. Res. Develop.* 4(5):518–524.
- [34] Li J, Tang P (2022) Auto-bidding equilibrium in ROI-constrained online advertising markets. Preprint, submitted October 12, <https://arxiv.org/abs/2210.06107>.
- [35] Mehta A (2013) Online matching and ad allocation. *Found. Trends Theor. Comput. Sci.* 8(4):265–368.
- [36] Mehta A, Saberi A, Vazirani U, Vazirani V (2007) Adwords and generalized online matching. *J. ACM* 54(5):22.
- [37] Othman A, Papadimitriou C, Rubinstein A (2016) The complexity of fairness through equilibrium. *ACM Trans. Econom. Comput.* 4(4):1–19.
- [38] Papadimitriou CH (1994) On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. System Sci.* 48(3):498–532.
- [39] Rosen A, Segev G, Shahaf I (2021) Can PPAD hardness be based on standard cryptographic assumptions? *J. Cryptology* 34(1):8.
- [40] Roughgarden T (2020) Complexity theory, game theory, and economics: The Barbados lectures. *Found. Trends Theor. Comput. Sci.* 14(3–4):222–407.
- [41] Vazirani VV, Yannakakis M (2011) Market equilibrium under separable, piecewise-linear, concave utilities. *J. ACM* 58(3):1–25.