Complexity of Equilibria in First-Price Auctions under General Tie-Breaking Rules

Xi Chen
Columbia University
New York, United States
xichen@cs.columbia.edu

Binghui Peng
Columbia University
New York, United States
bp2601@columbia.edu

ABSTRACT

We study the complexity of finding an approximate (pure) Bayesian Nash equilibrium in a first-price auction with common priors when the tie-breaking rule is part of the input. We show that the problem is PPAD-complete even when the tie-breaking rule is trilateral (i.e., it specifies item allocations when no more than three bidders are in tie, and adopts the uniform tie-breaking rule otherwise). This is the first hardness result for equilibrium computation in first-price auctions with common priors. On the positive side, we give a PTAS for the problem under the uniform tie-breaking rule.

CCS CONCEPTS

• **Theory of computation → Exact and approximate computation of equilibria**.

KEYWORDS

Equilibrium computation, First-price auction

ACM Reference Format:
Xi Chen and Binghui Peng. 2023. Complexity of Equilibria in First-Price Auctions under General Tie-Breaking Rules. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC '23), June 20–23, 2023, Orlando, FL, USA*. ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3564246.3585195

1 INTRODUCTION

First-price auction is arguably the most commonly used auction format in practice [15, 46, 54], in which the highest bidder wins the item and pays her bid. First-price auction and its variants have been widely used in online ad auctions: when a user visits a platform, an auction is run among interested advertisers to determine the ad to be displayed to the user. Despite of its simplicity, first-price auctions are *not* incentive compatible — it is the most well-known example in auction theory that does not admit a truthful strategy. This has led to significant effort in economics [2, 3, 32–35, 37, 38, 44] and more recently, in computer science [6, 24, 55], to understand equilibria of first-price auctions.

In this paper we study the computational complexity of finding a Bayesian Nash equilibrium in a first-price auction. We consider the following *independent common prior* setting. There is one single item to sell, and n bidders are interested in it. Each bidder has a continuous value distribution $\mathcal{D}_i$ supported over $[0, 1]$. The joint value distribution $\mathcal{D}$ is the product of $\mathcal{D}_i$'s. While $\mathcal{D}_i$'s are public, each bidder i has a private value v_i for the item drawn from $\mathcal{D}_i$. Each bidder chooses a bidding strategy, which maps her private value to a bid from a *discrete* bid space $\mathcal{B} = \{b_0, b_1 \ldots, b_m\}$. A Bayesian Nash equilibrium is a tuple of bidding strategies, one for each bidder, such that every bidder gets a best response to other bidders' strategies (see formal definition in Section 2, including the two conditions that bidding strategies need to satisfy: no overbidding and monotonicity).

This game between bidders, however, is not fully specified without a tie-breaking rule: how the item is allocated when more than one bidder have the highest bid. A variety of tie-breaking rules have been considered in the literature. The uniform tie-breaking rule, where the item is allocated to one of the winners uniformly at random, has been the most common offset. The other commonly used tie-breaking rule is to perform an additional round of Vickrey auction to ensure the existence of equilibria when the bidding space is continuous [22, 32]. A recent line of works [26, 28, 29, 51] used monopoly tie-breaking rules that always give the item to one player when establishing worst-case price-of-anarchy (POA) bounds.

To accommodate tie-breaking rules in the problem, we consider the setting where the auctioneer specifies a tie-breaking rule Γ to be used in the auction (as part of the input). Γ maps each $W \subseteq [n]$ as the set of winners to a distribution $\Gamma(W)$ over W as the allocation of the item to bidders in W. While a general tie-breaking rule takes exponentially many entries to describe, our PPAD-hardness result is built upon the succinct family of so-called *trilateral* tie-breaking rules: such a tie-breaking rule Γ specifies item allocations when no more than three bidders are in tie, and follows the uniform tie-breaking rule otherwise (when more than three bidders are in tie). The hardness result rules out the possibility of an efficient algorithm for finding a Bayesian Nash equilibrium in a first-price auction when the tie-breaking rule is given as part of the input (unless PPAD is in P). We compliment our hardness result with a polynomial time approximation scheme (PTAS) for finding a constant-approximate Bayesian Nash equilibrium under the uniform tie-breaking rule.

1.1 Our Results

Our main hardness result shows that the problem of finding an ϵ-approximate Bayesian Nash equilibrium in a first-price auction is PPAD-complete with trilateral tie-breaking.

Theorem 1.1 (Computational hardness). *It is PPAD-complete to find an ϵ-approximate Bayesian Nash equilibrium in a first-price auction under a trilateral tie-breaking rule for $\epsilon = 1/\operatorname{poly}(n)$.*

STOC '23, June 20–23, 2023, Orlando, FL, USA

ACM ISBN 978-1-4503-9913-5/23/06…$15.00
https://doi.org/10.1145/3564246.3585195

It is worth pointing out that the hardness above holds even when (1) the bid space $\mathcal{B}$ has size 3; and (2) the density function of each $\mathcal{D}_i$ is a piecewise-constant function with no more than four nonzero pieces.

On the positive side, we obtain a PTAS for finding a Bayesian Nash equilibrium in a first-price auction under the uniform tie-breaking rule:

Theorem 1.2 (PTAS under uniform tie-breaking). *For any $\epsilon > 0$, $n, m \geq 2$, there is an algorithm that finds an ϵ-approximate Bayesian Nash equilibrium using $O(n^4 \cdot g(1/\epsilon))$ time under the uniform tie-breaking rule.*

Our algorithm works as long as it has oracle access to the CDF of each value distribution $\mathcal{D}_i$.

1.2 Related Work

First-price auction. The study of first-price auction dates back to the seminal work of Vickrey [54] in 1960s. Despite its extremely simple form and a wide range of applications, the incentive has been a central issue and it is perhaps the most well known mechanism that does not admit a truthful strategy. A long line of works in the economic literature [2, 3, 6, 32–38, 43–45] devote to characterizing the existence, uniqueness and closed-form expression of a pure Bayesian Nash equilibrium (or BNE). However, the BNE of first-price auction is only well-understood in a few special cases, including when the players have symmetric valuation distributions [6], when all players have probability density function bounded above 0 and atomic probability mass at the lowest points [34], when there are only two bidders with uniform valuation distributions [31] or when the players have discrete value and continuous bidding space and the tie-breaking is performed with an extra round of Vickrey (second-price) auction [55].

A formal study on the computational complexity of equilibria in a first-price auction has been raised by the recent work of [24], which is most closest to us. [24] examines the computation complexity under a *subjective prior*, that is, each bidder has a different belief of other's valuation distribution. They prove the PPAD-completeness and the FIXP-completeness of finding an ϵ-BNE (for some constant $\epsilon > 0$) and an exact BNE, under the uniform tie-breaking rule. As we shall explain soon, the techniques to obtain their results are quite different from us. It is worth noting that most aforementioned literature are on the common prior setup, and [24] also leaves an open question of characterizing the computational complexity of ϵ-BNE under the standard setting of independent common prior. [24] also provides a polynomial time algorithm for finding a high precision BNE for *constant* number of players and bids, when the input distribution are piecewise polynomial. Their approach is based on polynomial system solvers and thus different from us. The work of [4] studies the Bayesian combinatorial auctions, where there are multi-items to sell for multiple bidders. They prove the complexity of Bayesian Nash equilibrium is at least PP-hard (a complexity class between the polynomial hierarchy and PSPACE), the model is quite different, because the agents' valuation could be much more complex, defining over subsets of items.

Other aspects of first-price auction have also been studied in the literature, including the price of anarchy/stability [23, 27–29, 51, 52] and parameter estimation [14, 25].

Equilibrium computation. The complexity class of PPAD (Polynomial Parity Arguments on Directed graphs) was first introduced by Papadimitriou [42] to capture one particular genre of total search functions. The seminal work [9, 16] established the PPAD-hardness of normal-form games. The hardness of approximation was settled by subsequent work [21, 48, 49] in the past few years. A broad range of problems have been proved to be PPAD-hard, and notable examples including equilibrium computation in special but important class of games (win-or-lose game [1, 13], anonymous game [10], constant rank game [39], graphical game [41]), market equilibrium (Arrow-Debreu market [8, 12, 53], non-monotone market [11, 50], Hylland-Zeckhauser scheme [7]), fair division [5, 40], min-max optimization [19] and reinforcement learning [17, 30].

The PTAS is known for anonymous game [18], which is closely related to our work. The [18] presented a $n^{g(m,1/\epsilon)} \cdot U$ algorithm for m-action n-player anonymous games for some exponential function g. Here U denotes the number of bits to represent a payoff value in the game. Instead, our algorithm finds an ϵ-BNE of first-price auction with running time $n^4 \cdot g(1/\epsilon)$, which does not depend on the size of bidding space and the bit-size of the representation of the distributions. It crucially utilizes the structure of first-price auction in the rounding and searching step, and could have a broader application in auction theory.

1.3 Technical Overview

The challenge of obtaining the PPAD-hardness arises from two folds. First, the utility function does not admit a closed-form expression, in terms of other player's strategy. It depends on an exponential number of possible bidding profiles and is computed only via a dynamic programming approach. Second, the game structure is highly symmetric under the (independent) common prior. In a first-price auction, the allocation is determined by the entire bidding profile, and each player faces "almost" the same set of profile. From this perspective, it is more like an anonymous game. Perhaps even worse, in an anonymous game, the utility function of each player is different and could be designed for the sake of reduction. While in a first-price auction, the utility function of each player is the same, and depends only on the allocation probability. Of course, the general (non-uniform) tie-breaking rule as well as the different valuation distributions could be used for breaking the symmetry. We note the above challenges are unique to the common prior setting. In a sharp contrast, in the subjective prior setting [24], the players' subjective belief could be different. A player could presume most other players have zero value and submit zero bid, hence, the game is *local* and non-symmetric.

To resolve the above challenges, our key ideas are (1) linearizing the allocation probability and expanding a first order approximation of the utility function; and (2) carefully incorporating a (simple) general tie-breaking rule to break the symmetry.

Technical highlight: Linearizing "everything". Given a strategy profile s, the distribution over the entire bidding profile (and therefore the allocation probability, the utility, the best response) could be complicated to compute, especially when multiple players submit the highest bid. To circumvent this issue, we assign a large probability $(1 - \delta)$ around value 0 for all players, for some polynomially small $\delta > 0$.[1] By doing this, the probability that a player bids nonzero is small, so one can ignore higher order term. Concretely, let $p_{i,j}$ be the probability that player i gets the item when bidding b_j given that the other players have strategy s_{-i}, and let $\Gamma_i(b_j, s_{-i})$ be the allocation for player i of bidding b_j given other player's strategy s_{-i}. The immediate advantage is that the allocation probability can be approximated as

$$\Gamma_i(b_j, s_{-i}) \approx (1 - \sum_{i' \in [n]\setminus\{i\}} \sum_{j' > j} p_{i',j'}) + \sum_{i' \in [n]\setminus\{i\}} \Sigma_{i,i'} \cdot p_{i',j} \quad (1)$$

under a *bilateral* tie-breaking rule. Here $\Sigma \in [0,1]^{n\times n}$ specifies the allocation when there is a tie between a pair of players (i_1, i_2) and satisfies $\Sigma + \Sigma^\top = (J - I)$. At this stage, it is tempting to use $p_{i,j}$ to encode variables of a generalized circuit problem and the choice of best response to encode constraints. In our final construction, we only need three bids $0 = b_0 < b_1 < b_2$ and the variables are encoded by the jump point τ_i between b_1, b_2 (i.e., when player i bids b_2 instead of b_1), which has the closed-form expression of

$$\tau_i = b_2 + \frac{\Gamma_i(b_1, s_{-i}) \cdot (b_2 - b_1)}{\Gamma_i(b_2, s_{-i}) - \Gamma_i(b_1, s_{-i})}. \quad (2)$$

Even after the linearization step of Eq. (1), the above expression is still quite formidable to handle. Our next idea is to restrict the jumping point in a small interval between $(b_2, 1)$, and assign only a small total probability mass of $\beta\delta$ over the interval, here β is another polynomially small value. There is a (fixed) probability mass of δ around b_2 and 1. One can further perform a first order approximation to Eq. (2), and again linearize the jumping point expression.

Incorporating the tie-breaking rule. Abstracting away some construction details, the above construction reduces the first-price auction from a fix point problem, obeys the following form

$$\vec{p} = f(G\vec{p}) \quad \text{where} \quad G = 2\Sigma^A - J + I, \quad (3)$$

where $\vec{p}$ is the probability of bidding b_1 (inside the small interval), $f = (f_1, \dots, f_n)$ is operated coordinate-wise over $G\vec{p}$, f_i is a monotone function maps from $[a_i, b_i]$ (some fixed interval) to $[0,1]$, J is the all 1 matrix and I is the identity matrix. The fixed point problem is fairly general and subsumes the generalized circuit problem, *if* Σ^A is an arbitrary matrix in $[0,1]^{n\times n}$. Unfortunately, it is not true due to the constraint of $\Sigma^A + (\Sigma^A)^\top = (J - I)$. We resolve the issue by adding an extra pivot player. The pivot player is guaranteed to bid $b_0 = 0$ and b_2 with equal probability of 1/2. From a high level, the pivot player splits the equilibrium computation into two cases, the case when it bids b_0 is similar, while the case of bidding b_2 introduces another tie-breaking matrix $\Sigma^B \in [0,1]^{n\times n}$ among the original players in $[n]$ (hence it becomes a trilateral rule). It transforms the fix point problem (i.e., Eq. (3)) to a more convenient form

$$\vec{p} = f(G'\vec{p}) \quad \text{where} \quad G' = 2\Sigma^A + \Sigma^B - J + I, \quad (4)$$

and one can construct gadgets to reduce from the generalized circuit problem. The last step is fairly common and details can be found in Section 3.

[1]This is the reason that our hardness result only applies for (inverse) polynomially small ϵ.

2 PRELIMINARY

Notation. We write $[n]$ to denote $\{1, 2, \dots, n\}$ and $[n_1 : n_2]$ to denote $\{n_1, n_1 + 1, \dots, n_2\}$. Let 1_i be an indicator vector – it equals the all 0 vector, except the i-th coordinate which equals 1. Let Δ_n contains all probability distribution over $[n]$. Given a vector $v \in \mathbb{R}^n$, and an index $i \in [n]$, v_i denotes the i-th entry of v while v_{-i} denotes $(v_1, v_2, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$, i.e., all entries except the i-th coordinate. We write $x = y \pm \epsilon$ if $x \in [y - \epsilon, y + \epsilon]$. Let $J_n \in \mathbb{R}^{n\times n}$ be the $n \times n$ all-1 matrix and I_n be the $n \times n$ identity matrix.

2.1 Model

In a Bayesian first-price auction (FPA), there is one single item to sell and it is specified by a tuple $(\mathcal{N}, \mathcal{B}, \mathcal{D}, \Gamma)$, where $\mathcal{N} = [n]$ is the set of players, $\mathcal{B}$ is the bid space, $\mathcal{D}$ is the value distribution and Γ is the tie-breaking rule. For each play $i \in \mathcal{N}$, it has a private value v_i of the item that is drawn from a (continuous) distribution $\mathcal{D}_i$ supported over $[0,1]$ (written as $v_i \sim \mathcal{D}_i$). We consider the standard *independent common prior* setting — the joint value distribution $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$ is the product distribution of $\{\mathcal{D}_i\}_{i\in[n]}$ and we assume the value profile $v = (v_1, \dots, v_n) \in [0,1]^n$ is drawn from $\mathcal{D}$. Let $\mathcal{B} = \{b_0, b_1, \dots, b_m\} \subset [0,1]$ be the bid space, where $0 = b_0 < b_1 < \cdots < b_m \leq 1$.

In a first-price (sealed-bid) auction, each bidder i submits a bid $\beta_i \in \mathcal{B}$ simultaneously to the seller. The seller assigns the item to the winning player i^* which submits the highest bid, and charges i^* a payment equals to its bid β_{i^*}.

Allocation and tie-breaking rule. When there are multiple players submitting the same highest bid, the seller assigns and charges the item to one of those winning players, following a pre-described *tie-breaking* rule Γ. A tie breaking rule $\Gamma : \{0,1\}^n \to \Delta_n$ maps a set of winning players $W \subseteq [n]$ to an allocation profile $\Gamma(W) \in \Delta_n$ supported on W that specifies the winning probability of each player $i \in W$ as $\Gamma_i(W)$. Formally, given a bidding profile $\beta \in \mathcal{B}^n$, the set of winning players $W(\beta)$ are those who submit the highest bids

$$W(\beta) = \left\{ i \in [n] : \beta_i = \max_{j\in[n]} \beta_j \right\}.$$

The tie breaking rule $\Gamma(W(\beta)) \in \Delta_n$ specifies the winning probability of each player in $W(\beta)$ and $\Gamma_i(W(\beta))$ is the probability that the bidder i obtains the item under the bidding profile β. The tie-breaking rule needs to satisfy (1) $\Gamma_i(W(b)) > 0$ only if $i \in W(\beta)$, i.e., the item is assigned only to players with the highest bid; and (2) $\sum_{i\in[n]} \Gamma_i(W(\beta)) = 1$, i.e., the total allocation is 1. When there is no confusion, we also abbreviate $\Gamma(\beta) = \Gamma(W(\beta))$.

It is known that the tie-breaking rule plays a subtle yet critical rule on the equilibrium of Bayesian FPA. Our hardness result is built upon the *trilateral* tie-breaking rule, a simple generalization of the commonly used uniform tie-breaking method.

Definition 2.1 (Trilateral tie-breaking). *A* trilateral *tie-breaking rule Γ is specified by the following tuples of nonnegative numbers*

$$\left(w_{i,j} : 1 \leq i < j \leq n\right) \quad \textit{and} \quad \left(\sigma^{(1)}_{i,j,k}, \sigma^{(2)}_{i,j,k} : 1 \leq i < j < k \leq n\right)$$

such that $w_{i,j} \le 1$ and $\sigma^{(1)}_{i,j,k} + \sigma^{(2)}_{i,j,k} \le 1$. Given a bidding profile $\beta \in \mathcal{B}^n$ and the winning set $W(\beta)$, the item is distributed according to Γ as follows

(1) *If $W(\beta) = \{i\}$ for some $i \in [n]$, then $\Gamma_i(\beta) = 1$;*
(2) *If $W(\beta) = \{i, j\}$ for some $1 \le i < j \le n$, then $\Gamma_i(\beta) = w_{i,j}$ and $\Gamma_j(\beta) = 1 - w_{i,j}$;*
(3) *If $W(\beta) = \{i, j, k\}$ for some $1 \le i < j < k \le n$, then $\Gamma_i(\beta) = \sigma^{(1)}_{i,j,k}$, $\Gamma_j(\beta) = \sigma^{(2)}_{i,j,k}$ and $\Gamma_k(\beta) = 1 - \sigma^{(1)}_{i,j,k} - \sigma^{(2)}_{i,j,k}$; and*
(4) *When $|W(\beta)| \ge 4$, the item is evenly distributed among players in $W(\beta)$.* [2]

Equilibrium and strategy. Given a tie-breaking rule Γ and a bidding profile $\beta = (\beta_1, \dots, \beta_n)$, the *ex-post* utility of a bidder i is given by

$$u_i(v_i; \beta_i, \beta_{-i}) = (v_i - \beta_i) \cdot \Gamma_i(\beta).$$

A strategy $s_i : [0, 1] \to \mathcal{B}$ of player i is a map from her (private) value v_i to a bid $s(v_i) \in \mathcal{B}$, with the following two properties:

- **No overbidding**. A player never submits a bid larger than her private value, i.e., $s_i(v_i) \le v_i$ for all $v_i \in [0, 1]$.
- **Monotonicity.** s_i is a non-decreasing function.

These are common assumptions in the literature of first-price auction [24, 34, 38] and they rule out spurious equilibria in Bayesian auctions [4]. Due to the monotonicity assumption, one can write a strategy s_i as m thresholds $0 \le \tau_{i,1} \le \cdots \le \tau_{i,m} \le 1$, where the player i bids b_j in the interval $(\tau_{i,j}, \tau_{i,j+1}]$[3]. Here we set by default $\tau_{i,0} = 0$ and $\tau_{i,m+1} = 1$.

The ϵ-approximate Bayesian Nash equilibrium (ϵ-approximate BNE) of FPA is defined as follow.

Definition 2.2 (ϵ-approximate Bayesian Nash equilibrium)**.** *Let $n, m \ge 2$. Given a first-price auction $(\mathcal{N}, \mathcal{B}, \mathcal{D}, \Gamma)$, a strategy profile $s = (s_1, \dots, s_n)$ is an ϵ-approximate Bayesian Nash equilibrium (ϵ-approximate BNE) if for any player $i \in [n]$, we have*

$$\mathop{\mathbb{E}}_{v\sim\mathcal{D}}\Big[u_i(v_i; s_i(v_i), s_{-i}(v_{-i}))\Big] \ge \mathop{\mathbb{E}}_{v\sim\mathcal{D}}\Big[u_i(v_i; \mathrm{bs}(v_i, s_{-i}), s_{-i}(v_{-i}))\Big] - \epsilon,$$

where $\mathrm{bs}(v_i, s_{-i}) \in \mathcal{B}$ is the best response of player i given other players' strategy s_{-i}, i.e.

$$\mathrm{bs}(v_i, s_{-i}) \in \arg\max_{b\in\mathcal{B}} \mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}\Big[u_i(v_i; b, s_{-i}(v_{-i}))\Big].$$

The existence and the PPAD membership of finding a $1/\mathrm{poly}(n)$-approximate BNE can be established via a similar approach of [24] (In particular, Theorem 4.1 and Theorem 4.4 of [24]), and we omit the standard proof here.

We shall also use another notion of equilibrium which is more convenient in our hardness reduction. The ϵ-approximately well-supported Bayesian Nash equilibrium (ϵ-BNE) is defined as

Definition 2.3 (ϵ-approximately well-supported Bayesian Nash equilibrium)**.** *Let $n, m \ge 2$. Given a first-price auction $(\mathcal{N}, \mathcal{B}, \mathcal{D}, \Gamma)$, a strategy profile $s = (s_1, \dots, s_n)$ is an ϵ-approximately well-supported Bayesian Nash equilibrium (ϵ-BNE) if for any player $i \in [n]$ and $v_i \in [0, 1]$, we have*

$$\begin{aligned}&\mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}\Big[u_i(v_i; s_i(v_i), s_{-i}(v_{-i}))\Big]\\ \ge\ &\mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}\Big[u_i(v_i; \mathrm{bs}(v_i, s_{-i}), s_{-i}(v_{-i}))\Big] - \epsilon.\end{aligned}$$

Remark 2.4. *We note the ϵ-approximate BNE is known also* ex-ante *approximate BNE, and the ϵ-BNE is known as* ex-interim *approximate BNE in some of the literature.*

The notion of ϵ-BNE and ϵ-approximate BNE can be reduced to each other in polynomial time, losing at most a polynomial factor of precision. It is clear that an ϵ-BNE is also an ϵ-approximate BNE. Lemma 2.5 states the other direction and the proof can be found at the full version of this paper.

Lemma 2.5. *Given a first-price auction $(\mathcal{N}, \mathcal{B}, \mathcal{D}, \Gamma)$ and an ϵ-approximate BNE s, there is a polynomial time algorithm that maps s to an ϵ'-BNE, where $\epsilon' = (2n + 10)\sqrt{\epsilon}$.*

3 PPAD-HARDNESS

Recall our main hardness result

Theorem 1.1 (Computational hardness). *It is PPAD-complete to find an ϵ-approximate Bayesian Nash equilibrium in a first-price auction under a trilateral tie-breaking rule for $\epsilon = 1/\mathrm{poly}(n)$.*

In the rest of section, we construct the hard instances of FPA in Section 3.1 and provide some basic properties in Section 3.2. We reduce from the ϵ-generalized-circuit problem in Section 3.3.

3.1 Construction of First-Price Auctions

It suffices to prove finding ϵ-BNE is hard for some $\epsilon = 1/\mathrm{poly}(n)$ due to Lemma 2.5. We will use the following three parameters in the construction:

$$\epsilon = \frac{1}{n^{40}}, \quad \delta = \frac{1}{n^{10}} \quad \text{and} \quad \beta = \frac{1}{n^4}.$$

We describe the bidding space $\mathcal{B}$, the valuation distribution $\mathcal{D}$ and the tie-breaking rule Γ.

Bidding space. The bidding space $\mathcal{B} = \{b_0, b_1, b_2\}$ contains 3 bids in total, where $b_0 = 0$,

$$b_1 = \frac{\delta^2}{n^4} \quad \text{and} \quad b_2 = \frac{\delta}{n^2}.$$

Valuation distribution. There are $n + 1$ players — n standard players indexed by $[n]$ and one pivot player $n + 1$. We will describe the value distribution $\mathcal{D}_i$ of player i by specifying its density function $p_i : [0, 1] \to \mathbb{R}^+$. The density function p_{n+1} of the pivot player is set as follows:

$$p_{n+1}(v) = \begin{cases} 1/(2\epsilon) & v \in [0, \epsilon] \\ 1/(2\epsilon) & v \in [1 - \epsilon, 1] \end{cases}.$$

In another word, $\mathcal{D}_{n+1}$ has 0.5 probability mass around 0 and 0.5 probability mass around 1.

[2] We note our hardness result actually holds regardless of the tie-breaking rule among more than 3 players (i.e., not necessarily uniform).

[3] If the valuation distribution contains a point mass, then the strategy might be randomized at the point mass.

The density function p_i of each standard player $i \in [n]$ is set as follows:

$$p_i(v) = \begin{cases} (1-(2+\beta)\delta)/\epsilon & v \in [0,\epsilon] \\ \delta/\epsilon & v \in [b_2-\epsilon, b_2] \\ \widetilde{p_i}(v) & v \in (b_2, 1-\epsilon) \\ \delta/\epsilon & v \in [1-\epsilon, 1] \end{cases}$$

where $\widetilde{p_i}(v)$ is defined over $(b_2, 1-\epsilon)$, satisfies $\int_{b_2}^{1-\epsilon} \widetilde{p_i}(v)\mathrm{d}v = \beta\delta$, but will be specified later in the reduction in Section 3.3. In short, a standard player i has most its probability mass around 0, δ mass around b_2, δ mass around 1 and $\beta\delta$ mass in $(b_2, 1-\epsilon)$ to be specified later.

Tie-breaking rule. We describe the trilateral tie-breaking rule Γ as follows. For any bidding profile β with $2 \le |W(\beta)| \le 3$, the tie-breaking rule depends on the presence of $n+1$ in $W(\beta)$:

- Suppose $n+1 \notin W(\beta)$. Then
  - If $|W(\beta)| = 2$, i.e., $W(b) = \{i,j\}$, the tie-breaking rule is given by a matrix $\Sigma^A \in [0,1]^{n\times n}$ such that player i obtains $\Sigma^A_{i,j}$ unit of the item and player j obtains $\Sigma^A_{j,i}$ unit. So the matrix Σ^A needs to satisfy $\Sigma^A + (\Sigma^A)^\top = (J_n - I_n)$. We will specify Σ^A in the reduction later but will guarantee that all of its off-diagonal entries lie in $[1/4, 3/4]$.
  - If $|W(\beta)| = 3$, then we use the uniform allocation.
- Suppose $n+1 \in W(\beta)$. Then
  - If $|W(\beta)| = 2$, then the item is fully allocated to the pivot player $n+1$.
  - If $|W(\beta)| = 3$, i.e., $W(b) = \{i,j,n+1\}$, then the tie breaking is given by a matrix $\Sigma_B \in [0,1]^{n\times n}$ such that player i obtains $\Sigma^B_{i,j}$ unit of the item, player j obtains $\Sigma^B_{j,i}$ unit and player $n+1$ obtains $1-\Sigma^B_{i,j}-\Sigma^B_{j,i}$ unit. So the matrix Σ^B needs to satisfy $\Sigma^B + (\Sigma^B)^\top \le J_n - I_n$, i.e., $\Sigma^B + (\Sigma^B)^\top$ is entrywise dominated by $J_n - I_n$.

3.2 Basic Properties

Let $s = (s_1, \dots, s_{n+1})$ be an ϵ-BNE of the instance. We prove a few properties of s in this subsection. Given s, for each player i we define $f_i : \mathcal{B} \to [0,1]$ and $F_i : \mathcal{B} \to [0,1]$ as follows:

$$f_i(b) = \Pr_{v_i\sim\mathcal{D}_i}[s_i(v_i) = b] \quad \text{and} \quad F_i(b) = \Pr_{v_i\sim\mathcal{D}_i}[s_i(v_i) \le b].$$

In the rest part of section, we abbreviate

$$\Gamma_i(b, s_{-i}) := \mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}[\Gamma_i(b, s_{-i}(v_{-i}))]$$

and

$$u_i(v_i; b, s_{-i}) := \mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}[u_i(v_i; b; s_{-i}(v_{-i}))]$$

when there is no confusion.

We start with the following lemma.

Lemma 3.1 (Separable bid)**.** *In any ϵ-BNE, the equilibrium strategy s satisfies*

- *For a standard player $i \in [n]$, its equilibrium strategy satisfies*
  - *when $v_i \in [0,\epsilon]$, $s_i(v_i) = b_0$;*
  - *when $v_i \in [b_2-\epsilon, b_2]$, $s_i(v_i) = b_1$; and*
  - *when $v_i \in [1-\epsilon, 1]$, $s_i(v_i) = b_2$.*
- *For the pivot player, its equilibrium strategy satisfies*
  - *when $v_{n+1} \in [0,\epsilon]$, $s_{n+1}(v_{n+1}) = b_0$; and*
  - *when $v_{n+1} \in [1-\epsilon, 1]$, $s_{n+1}(v_{n+1}) = b_2$.*

Proof. The claim of $s_i(\epsilon) = 0$ holds trivially for all $i \in [n+1]$ due to the no-overbidding assumption. A standard player i chooses between b_0 and b_1 for $v_i \in [b_2-\epsilon, b_2]$. The allocation probability $\Gamma_i(b_0, s_{-i})$ of bidding b_0 satisfies $\Gamma_i(b_0, s_{-i}) \le \frac{1}{n+1}$, hence the utility of bidding $b_0 = 0$ is at most

$$u_i(v_i; b_0, s_{-i}) = (v_i - b_0)\cdot\Gamma_i(b_0, s_{-i}) \le \frac{1}{n} b_2.$$

The allocation probability of bidding b_1 is at least

$$\Gamma_i(b_1, s_{-i}) \ge \prod_{i\in[n+1]} f_i(b_0) \ge (1-(2+\beta)\delta)^n \cdot \frac{1}{2} \ge \frac{1}{3}$$

hence the utility of bidding b_1 is at least

$$\begin{aligned} u_i(v_i; b_1, s_{-i}) &= (v_i - b_1)\cdot\Gamma_i(b_1, s_{-i}) \\ &\ge (b_2 - \epsilon - b_1)\cdot\frac{1}{3} \\ &> u_i(v; b_0, s_{-i}) + \epsilon. \end{aligned}$$

Finally, we analyse the equilibrium strategy around $v \in [1-\epsilon, 1]$ for all $n+1$ players. Via an analysis similar to the above argument, it is clear that both standard players and the pivot player would choose between b_1 and b_2. For a standard player $i \in [n]$, the allocation probability of bidding b_2 satisfies

$$\begin{aligned} &\Gamma_i(b_2, s_{-i}) \\ &\ge F_{n+1}(b_1)\cdot\prod_{j\in[n]\setminus[i]} F_j(b_1) \\ &\ge F_{n+1}(b_1)\cdot\left(\prod_{j\in[n]\setminus[i]} f_j(b_0) + \sum_{j\in[n]\setminus[i]} f_j(b_1)\prod_{r\in[n]\setminus\{i,j\}} f_r(b_0)\right) \\ &\ge \frac{1}{2}\left(\prod_{j\in[n]\setminus[i]} f_j(b_0) + \sum_{j\in[n]\setminus[i]} f_j(b_1)\prod_{r\in[n]\setminus\{i,j\}} f_r(b_0)\right) \\ &\quad + \frac{1}{2} f_{n+1}(b_1), \end{aligned} \tag{5}$$

where the last step holds as $f_{n+1}(b_0) = \frac{1}{2}$ and

$$\prod_{j\in[n]\setminus[i]} f_j(b_0) \ge (1-(2+\beta)\delta)^{n-1} \ge \frac{1}{2}. \tag{6}$$

The allocation probability of bidding b_1 satisfies

$$\begin{aligned} &\Gamma_i(b_1, s_{-i}) \\ &\le f_{n+1}(b_0)\cdot\left(\prod_{j\in[n]\setminus[i]} f_j(b_0) + \Phi + 9n^2\delta^2\right) \\ &\quad + f_{n+1}(b_1)\cdot\sum_{j\in[n]\setminus\{i\}} f_j(b_1) \\ &\le \frac{1}{2}\left(\prod_{j\in[n]\setminus[i]} f_j(b_0) + \sum_{j\in[n]\setminus[i]} f_j(b_1)\cdot\Sigma^A_{i,j}\prod_{r\in[n]\setminus\{i,j\}} f_r(b_0)\right) \\ &\quad + f_{n+1}(b_1)\cdot 3n\delta + 9n^2\delta^2. \end{aligned} \tag{7}$$

where

$$\Phi := \sum_{j\in[n]\setminus[i]} f_j(b_1)\cdot\Sigma^A_{i,j}\prod_{r\in[n]\setminus\{i,j\}} f_r(b_0)$$

Here the first step holds since (1) when the pivot player bids b_0, player i obtains $\Sigma^A_{i,j}$ unit of item when (only) player j bids b_1, and the probability of at least two players bidding b_1 is bounded as

$$\sum_{j,r\in[n]\setminus\{i\}} \Pr[s_r(v_r)=b_1\wedge s_j(v_j)=b_1]\le n^2\cdot 9\delta^2,$$

(2) when the pivot player bids b_1, the player i obtains the item only if there exists at least one other standard player j bids b_1 as the tie breaking rule assigns the item fully to player $n+1$ when there are only two winners. The second step follows from $f_j(b_1)\le 3\delta$ and that the pivot player does not bid b_0 in $[1-\epsilon,1]$ so $f_{n+1}(b_0)=1/2$.

Subtracting Eq. (7) and Eq. (5), one obtains

$$\begin{aligned}
&\Gamma_i(b_2,s-i)-\Gamma_i(b_1,s_{-i})\\
&\ge \frac{1}{2}\sum_{j\in[n]\setminus[i]} f_j(b_1)\cdot(1-\Sigma^A_{i,j})\prod_{r\in[n]\setminus\{i,j\}} f_r(b_0)]\\
&\quad+\frac{1}{2}f_{n+1}(b_1)-3n\delta f_{n+1}(b_1)-9n^2\delta^2\\
&\ge \frac{1}{2}\cdot(n-1)\cdot\delta\cdot\frac{1}{4}\cdot\frac{1}{2}+\frac{1}{2}f_{n+1}(b_1)-3n\delta f_{n+1}(b_1)-9n^2\delta^2\\
&\ge\frac{n\delta}{32}.
\end{aligned}\tag{8}$$

The second step holds due to $f_j(b_1)\ge\delta,\Sigma^A_{i,j}\in[1/4,3/4]$ and Eq (6). Hence we claim player i prefers b_2 than b_1 at value $v_i\in[1-\epsilon,1]$, since

$$\begin{aligned}
&u_i(v_i;b_2,s_{-i})-u_i(v_i;b_1,s_{-i})\\
&=(v_i-b_2)\cdot\Gamma_i(b_2,s_{-i})-(v_i-b_1)\cdot\Gamma_i(b_1,s_{-i})\\
&\ge v_i\cdot(\Gamma_i(b_2,s_{-i})-\Gamma_i(b_1,s_{-i}))-b_2\\
&\ge(1-\epsilon)\cdot\frac{n\delta}{32}-b_2>\epsilon.
\end{aligned}$$

Finally, for the pivot player $n+1$, the allocation probability of bidding b_1 satisfies

$$\Gamma_{n+1}(b_1,s_{-i})\le\prod_{i\in[n]}F_i(b_1)\tag{9}$$

and the allocation probability of b_2 satisfies

$$\begin{aligned}
\Gamma_{n+1}(b_2,s_{-i})&\ge\prod_{i\in[n]}F_i(b_1)+\sum_{i\in[n]}f_i(b_2)\prod_{j\in[n]\setminus\{i\}}F_i(b_1)\\
&\ge\prod_{i\in[n]}F_i(b_1)+\frac{n\delta}{2}.
\end{aligned}\tag{10}$$

The first step holds since the tie-breaking rule favors player $n+1$ when at most one player in $[n]$ bids b_2, the second step holds due to $f_i(b_2)\ge\delta$ and Eq. (6). Hence, at any $v_{n+1}\in[1-\epsilon,1]$ we have

$$\begin{aligned}
&u_{n+1}(v_{n+1};b_2,s_{-i})-u_{n+1}(v;b_1,s_{-i})\\
&=(v_{n+1}-b_2)\cdot\Gamma_{n+1}(b_2,s_{-i})-(v_{n+1}-b_1)\cdot\Gamma_{n+1}(b_1,s_{-i})\\
&\ge v_{n+1}\cdot(\Gamma_{n+1}(b_2,s_{-i})-\Gamma_{n+1}(b_1,s_{-i}))-b_2\\
&\ge(1-\epsilon)\cdot\frac{n\delta}{2}-b_2>\epsilon.
\end{aligned}$$

We conclude the proof of the lemma here. □

Lemma 3.1 confirms that b_0,b_1,b_2 would appear in an ϵ-BNE profile for every player $i\in[n]$. It still remains to determine at which value point a standard player $i\in[n]$ jumps from b_1 to b_2 in s_i. Let $\tau_i\in(b_2,1-\epsilon)$ be the jumping point from b_1 to b_2 of a standard player i. The following formula is convenient to use.

Lemma 3.2 (Jumping point formula). *The jumping point τ_i of a standard player $i\in[n]$ satisfies*

$$\tau_i=b_2+\frac{\Gamma_i(b_1,s_{-i})\cdot(b_2-b_1)}{\Gamma_i(b_2,s_{-i})-\Gamma_i(b_1,s_{-i})}\pm O\left(\frac{\epsilon}{\delta}\right).$$

Proof. At any value point $v\in[0,1]$, recall the utility of bidding b_1 equals

$$u_i(v_i;b_1,s_{-i})=(v_i-b_1)\Gamma_i(b_1,s_{-i})$$

and the utility of bidding b_2 equals

$$u_i(v_i;b_2,s_{-i})=(v_i-b_2)\Gamma_i(b_2,s_{-i}).$$

Solving for $u_i(\tau_i,b_1,s_{-i})=u_i(\tau_i,b_2,s_{-i})\pm\epsilon$, one obtains

$$\begin{aligned}
\tau_i&=\frac{\Gamma_i(b_2,s_{-i})b_2-\Gamma_i(b_1,s_{-i})b_1\pm\epsilon}{\Gamma_i(b_2,s_{-i})-\Gamma_i(b_1,s_{-i})}\\
&=b_2+\frac{\Gamma_i(b_1,s_{-i})\cdot(b_2-b_1)\pm\epsilon}{\Gamma_i(b_2,s_{-i})-\Gamma_i(b_1,s_{-i})}\\
&=b_2+\frac{\Gamma_i(b_1,s_{-i})\cdot(b_2-b_1)}{\Gamma_i(b_2,s_{-i})-\Gamma_i(b_1,s_{-i})}\pm O\left(\frac{\epsilon}{\delta}\right).
\end{aligned}$$

The last step follows from Eq. (8), and this finishes the proof of the lemma. □

Let $x_i\in[0,\beta\delta]$ be the probability mass over the interval (b_2,τ_i), i.e., $x_i=\int_{b_2}^{\tau_i}p_i(v)\mathrm{d}v$, which we will refer to the jumping probability of s_i. We state a few facts that will be used repeatedly.

Lemma 3.3 (Basic facts).

- *For any standard player $i\in[n]$, we have*

$$\prod_{j\in[n]\setminus\{i\}}f_j(b_0)=1-(n-1)(2+\beta)\delta\pm O(n^2\delta^2)$$

and

$$\prod_{j\in[n]\setminus\{i\}}F_j(b_1)=1-(n-1)(\beta+1)\delta+\sum_{j\in[n]\setminus\{i\}}x_j\pm O(n^2\delta^2).$$

- *For any $e\in\{1,2\}$, we have*

$$\sum_{i\neq j\in[n]}\Pr\left[s_i(v_i)=b_e\wedge s_j(v_j)=b_e\right]=O(n^2\delta^2).$$

Proof. For the first claim, we have

$$\begin{aligned}
\prod_{j\in[n]\setminus\{i\}}f_j(b_0)&=(1-(2+\beta)\delta)^{n-1}\\
&=1-(n-1)(2+\beta)\delta\pm O(n^2\delta^2)
\end{aligned}$$

due to the choice of δ. Similarly we have (using $x_i\in[0,\beta\delta]$

$$\begin{aligned}
&\prod_{j\in[n]\setminus\{i\}}F_j(b_1)\\
&=\prod_{j\in[n]\setminus\{i\}}(1-(1+\beta)\delta+x_i)\\
&=1-(n-1)(1+\beta)\delta+\sum_{j\in[n]\setminus\{i\}}x_j\pm O(n^2\delta^2).
\end{aligned}$$

For the second claim, for any $e \in \{1,2\}$, we have

$$\sum_{i \neq j \in [n]} \Pr\left[s_i(v_i) = b_e \wedge s_j(v_j) = b_e\right] \leq \sum_{i \neq j \in [n]} (1+\beta)\delta^2 = O(n^2\delta^2).$$

We conclude the proof here. □

The key step is to determine the jumping point, where we use approximation.

Lemma 3.4 (Jumping point). *The jumping point τ_i of a standard player $i \in [n]$ satisfies*

$$\tau_i = \left(\frac{\Delta_{i,1} - \Delta_{i,2}}{\Delta_{i,1}^2} \pm O\left(\frac{\beta^2}{\Delta_{i,1}}\right)\right) \cdot b_2$$

where

$$\Delta_{i,1} := (n-1)\delta + \sum_{j \in [n]\setminus\{i\}} \left(\beta\delta \cdot \Sigma^A_{i,j} + (1+\beta)\delta \cdot \Sigma^B_{i,j}\right) \in \left[(n-1)\delta, 2n\delta\right]$$

and

$$\Delta_{i,2} := \sum_{j \in [n]\setminus\{i\}} \left(1 - 2\Sigma^A_{i,j} - \Sigma^B_{i,j}\right) x_j \in \left[-2n\beta\delta, n\beta\delta\right].$$

Proof. For any standard player $i \in [n]$, we compute when it jumps from b_1 to b_2 using the formula in Lemma 3.2. To do so, we first compute $\Gamma_i(b_1, s_{-i})$ and $\Gamma_i(b_2, s_{-i})$.

$$\begin{aligned}
&\Gamma_i(b_1, s_{-i}) \\
&= f_{n+1}(b_0) \cdot \left(\prod_{j \in [n]\setminus\{i\}} f_j(b_0) + \Psi \pm O(n^2\delta^2)\right) \\
&= \frac{1}{2}\left(1 - (n-1)(2+\beta)\delta + \sum_{j \in [n]\setminus\{i\}} (\delta + x_j)\Sigma^A_{i,j}\right) \pm O(n^2\delta^2).
\end{aligned} \tag{11}$$

where

$$\Psi_1 := \sum_{j \in [n]\setminus\{i\}} f_j(b_1) \prod_{j \in [n]\setminus\{i,j\}} f_j(b_0) \cdot \Sigma^A_{i,j}$$

Here the first step follows from the tie-breaking rule and the second claim of Lemma 3.3, the second step follows from $f_j(b_1) = \delta + x_j$ and the first claim of Lemma 3.3.

The allocation probability of bidding b_2 obeys

$$\begin{aligned}
&\Gamma_i(b_2, s_{-i}) \\
&= f_{n+1}(b_0) \cdot \left(\prod_{j \in [n]\setminus\{i\}} F_j(b_1) + \Psi_2 \pm O(n^2\delta^2)\right) \\
&\quad + f_{n+1}(b_2) \cdot \left(\sum_{j \in [n]\setminus\{i\}} f_j(b_2) \prod_{k \in [n]\setminus\{i,j\}} F_k(b_1) \cdot \Sigma^B_{i,j} \pm O(n^2\delta^2)\right) \\
&= \frac{1}{2}\left(1 - (n-1)(1+\beta)\delta + \sum_{j \in [n]\setminus\{i\}} x_j + \sum_{j \in [n]\setminus\{i\}} (\delta + \beta\delta - x_j)\Sigma^A_{i,j}\right) \\
&\quad + \frac{1}{2} \sum_{j \in [n]\setminus\{i\}} (\delta + \beta\delta - x_j)\Sigma^B_{i,j} \pm O(n^2\delta^2).
\end{aligned}$$

where

$$\Psi_2 := \sum_{j \in [n]\setminus\{i\}} f_j(b_2) \prod_{j \in [n]\setminus\{i,j\}} F_j(b_1) \cdot \Sigma^A_{i,j}.$$

The first step uses the tie breaking rule and requires some explanations. In particular, (1) when the pivot player $n+1$ bids b_0, the player i obtains 1 unit of item when other players bid less than b_2, $\Sigma_{A,i,j}$ unit of item when only player j bids b_2; we also make use of the second claim of Lemma 3.3 to omit the other case; (2) when the player $n+1$ bids b_2, the player i obtains 0 unit of good when no other players bid b_0 and obtains $\Sigma_{B,i,j}$ unit of goods when one other player j bids b_2, and we omit other cases using Lemma 3.3. The second step follows from Lemma 3.3 and $f_j(b_2) = \delta + \beta\delta - x_j$.

Combining the above expression, we have

$$\begin{aligned}
&\Gamma_i(b_2, s_{-i}) - \Gamma_i(b_1, s_{-i}) \\
&= \frac{1}{2}(n-1)\delta + \frac{1}{2} \sum_{j \in [n]\setminus\{i\}} \left(\beta\delta \cdot \Sigma^A_{i,j} + (1+\beta)\delta \cdot \Sigma^B_{i,j}\right) \\
&\quad + \frac{1}{2} \sum_{j \in [n]\setminus\{i\}} \left(1 - 2\Sigma^A_{i,j} - \Sigma^B_{i,j}\right) x_j \pm O(n^2\delta^2).
\end{aligned} \tag{12}$$

Let $\Delta_{i,1}$ and $\Delta_{i,2}$ be defined as in the statement of the lemma. Note that $\Delta_{i,1}$ does not depend on $\{x_j\}_{j \neq i}$ while $\Delta_{i,2}$ depends on $\{x_j\}_{j \neq i}$. It is easy to see that

$$\Delta_{i,1} \in \left[(n-1)\delta, 2n\delta\right] \quad \text{and} \quad \Delta_{i,2} = \left[-2(n-1)\beta\delta, (n-1)\beta\delta\right]. \tag{13}$$

Finally we can compute the jumping point τ_i using Lemma 3.2 as follows:

$$\begin{aligned}
\tau_i &= b_2 + \frac{\Gamma_i(b_1, s_{-i}) \cdot (b_2 - b_1)}{\Gamma_i(b_2, s_{-i}) - \Gamma_i(b_1, s_{-i})} \pm O\left(\frac{\epsilon}{\delta}\right) \\
&= b_2 + \frac{1 \pm O(n\delta)}{\Delta_{i,1} + \Delta_{i,2}} \cdot b_2 \pm O\left(\frac{\epsilon}{\delta}\right) \\
&= \left(\frac{\Delta_{i,1} - \Delta_{i,2}}{\Delta_{i,1}^2} \pm O\left(\frac{\beta^2}{\Delta_{i,1}}\right)\right) \cdot b_2
\end{aligned}$$

The second step follows from Eq. (12), $\Gamma_i(b_1, \beta_i) = (1/2) \pm O(n\delta)$ (see Eq. (11)) and the choice of b_1, b_2. The last step follows from Eq. (13). □

3.3 Reduction from Generalized Circuit

Given $\alpha < \beta$, we write $\mathsf{T}_{[\alpha,\beta]} : \mathbb{R} \to [\alpha, \beta]$ to denote the truncation function with

$$\mathsf{T}_{[\alpha,\beta]}(x) = \min\left\{\max\{x, \alpha\}, \beta\right\}.$$

We recall the generalized circuit problem [9] and present a simplified version from [24].

Definition 3.5 ((Simplified) generalized circuit). *A generalized circuit is a tuple (V, G), where V is a set of nodes and G is a collection of gates. Each node $v \in V$ is associated with a gate G_v that falls into one of two types $\{G_{1-}, G_+\}$: If G_v is a G_+ gate, then it has two input nodes $v_1, v_2 \in V \setminus \{v\}$; if it is a G_{1-} gate then it takes one input node $v_1 \in V \setminus \{v\}$. Given $\kappa > 0$, a κ-approximation solution to (V, G) is an assignment $x \in [0,1]^V$ such that for every node v:*

- *If G_v is a G_+ gate and takes input nodes $v_1, v_2 \in V\setminus\{v\}$, then $x_v = \mathsf{T}_{[0,1]}(x_{v_1} + x_{v_2} \pm \kappa)$*

- *If G_v is a G_{1-} gate and takes an input node $v_1 \in V \backslash \{v\}$, then* $x_v = \mathsf{T}_{[0,1]}(1 - x_{v_1} \pm \kappa)$.

The generalized circuit problem is known to be PPAD-hard for constant κ.

Theorem 3.6 ([20, 47]). *There is a constant $\kappa > 0$ such that it is PPAD-hard to find an κ-approximate solution of a generalized circuit.*

We prove Theorem 1.1 via a reduction from the generalized circuit problem.

Given an instance of generalized circuit defined over nodes set V ($|V| = m$), we let $V_1 = [m_1]$ be the set of nodes with gate G_+ and $V_2 = [m_1+1 : m]$ be the set of nodes with gate G_{1-}. We construct an instance of first price auction with $n = m_1 + 2(m - m_1) = 2m - m_1$ standard players and one pivot player.

Let $\mathcal{N} = \mathcal{N}_1 \cup \mathcal{N}_2 \cup \mathcal{N}_3$ be the set of standard players, where $\mathcal{N}_1 = [m_1]$, $\mathcal{N}_2 = [m_1 + 1 : m]$ and $\mathcal{N}_3 = [m + 1 : 2m - m_1]$. From a high level, we use players in $\mathcal{N}_1$ to represent the set of nodes with G_+ gates, players in $\mathcal{N}_2$ to represent the set of nodes with G_{1-} gates. Players in $\mathcal{N}_3$ are used in constructing G_{1-}. We first specify the probability density $\widetilde{p}$ over interval $(b_2, 1-\epsilon)$ and the tie-breaking matrices Σ^A and Σ^B to complete the description of the FPA instance.

- For player i in $\mathcal{N}_1$ (i.e., $i \in [m_1]$), its valuation distribution $\widetilde{p}_i$ is uniform over the interval
$$\left[\frac{1}{\Delta_{i,1}} \cdot b_2, \left(\frac{1}{\Delta_{i,1}} + \frac{1}{10} \cdot \frac{\beta\delta}{\Delta_{i,1}^2}\right) \cdot b_2\right]$$
with a total probability mass of $\beta\delta$. Let $i(1), i(2) \in [m] = \mathcal{N}_1 \cup \mathcal{N}_2$ be the input nodes of the G_+ gates, we set $\Sigma^B_{i,i(1)} = \Sigma^B_{i,i(2)} = 1/10$.
- For player $m_1 + j \in \mathcal{N}_2$ (i.e., $j \in [m - m_1]$), its valuation distribution $\widetilde{p}_{m_1+j}$ is uniform over
$$\left[\left(\frac{1}{\Delta_{m_1+j,1}} - \frac{1}{10} \cdot \frac{\beta\delta}{\Delta_{m_1+j,1}^2}\right) \cdot b_2, \frac{1}{\Delta_{m_1+j,1}} \cdot b_2\right]$$
with a total probability mass of $\beta\delta$. We set $\Sigma^A_{m_1+j,m+j} = 9/20$.
- For player $m + j \in \mathcal{N}_3$ (i.e., $j \in [m - m_1]$), its valuation distribution $\widetilde{p}_{m+j}$ is uniform over
$$\left[\left(\frac{1}{\Delta_{m+j,1}} + \frac{1}{10} \cdot \frac{\beta\delta}{\Delta_{m+j,1}^2}\right) \cdot b_2, \left(\frac{1}{\Delta_{m+j,1}} + \frac{1}{5} \cdot \frac{\beta\delta}{\Delta_{m+j,1}^2}\right) \cdot b_2\right]$$
with a total probability mass of $\beta\delta$. We set $\Sigma^A_{m+j,m_1+j} = 11/20$. Let $j(1) \in [m]$ be the input node of G_{1-}, then set $\Sigma^B_{m+j,j(1)} = 1/5$.
- For any entry of Σ^A that has not been determined above, we set it to be $1/2$, and for any entry of Σ^B that has not been determined, we set it to be 0.

It is easy to verify that Σ^A and Σ^B satisfy the following properties as promised earlier: (1) the off-diagonal entries of Σ^A lie in $[1/4, 3/4]$; (2) $\Sigma^A + (\Sigma^A)^\top = J_n - I_n$; and (3) the off-diagonal entries of Σ^B belong to $[0, 1/2]$.

Letting $\kappa := n\beta = 1/n^3$, we prove that any ϵ-BNE of the first price auction gives an $O(\kappa)$-approximate solution to the generalized circuit. Indeed the following lemma shows that by taking $x_i' = x_i/(\beta\delta)$, where x_i is the jumping probability of s_i, we obtain an $O(\kappa)$-approximate solution $(x_1', \ldots, x_m')$ to the input generalized circuit. This finishes the proof of Theorem 1.1.

Lemma 3.7. *Given an ϵ-BNE of the first price auction, suppose $(x_1, \ldots, x_n) \in [0, \beta\delta]^n$ be the tuple of jumping probabilities, then we have*

- *For any $i \in [m_1]$, $x_i = \mathsf{T}_{[0,\beta\delta]}(x_{i(1)} + x_{i(2)} \pm O(\kappa\beta\delta))$*
- *For any $j \in [m - m_1]$, $x_{m_1+j} = \mathsf{T}_{[0,\beta\delta]}(\beta\delta - x_{j(1)} \pm O(\kappa\beta\delta))$*

Proof. For the first claim, for any $i \in [m_1]$, one has
$$\Delta_{i,2} = \sum_{r \in [n] \backslash \{i\}} \left(1 - 2\Sigma^A_{i,r} - \Sigma^B_{i,r}\right) x_r = -\frac{1}{10} x_{i(1)} - \frac{1}{10} x_{i(2)},$$
where the second step follows from $\Sigma^A_{i,r} = 1/2$ for all $r \in [n] \backslash \{i\}$, $\Sigma^B_{i,i(1)} = \Sigma^B_{i,i(2)} = 1/10$ and $\Sigma^B_{i,r} = 0$ for all other $r \in [n] \backslash \{i, i_1, i_2\}$. By Lemma 3.4, one has
$$\begin{aligned}\tau_i &= \left(\frac{\Delta_{i,1} - \Delta_{i,2}}{\Delta_{i,1}^2} \pm O\left(\frac{\beta^2}{\Delta_{i,1}}\right)\right) \cdot b_2 \\ &= \left(\frac{1}{\Delta_{i,1}} + \frac{1}{10} \cdot \frac{x_{i(1)} + x_{i(2)}}{\Delta_{i,1}^2} \pm O\left(\frac{\beta^2}{\Delta_{i,1}}\right)\right) \cdot b_2.\end{aligned}$$
Since $\mathcal{D}_i$ is uniform over $[\frac{1}{\Delta_{i,1}} \cdot b_2, (\frac{1}{\Delta_{i,1}} + \frac{1}{10} \cdot \frac{\beta\delta}{\Delta_{i,1}^2}) \cdot b_2]$ with probability mass $\beta\delta$, we have
$$x_i = \mathsf{T}_{[0,\beta\delta]}(x_{i(1)} + x_{i(2)} \pm O(\kappa\beta\delta)).$$

For the second claim, we first analyse the jumping probability of player $m_1 + j$. We have
$$\Delta_{m_1+j,2} = \sum_{r \in [n] \backslash \{m_1+j\}} \left(1 - 2\Sigma^A_{m_1+j,r} - \Sigma^B_{m_1+j,r}\right) x_r = \frac{1}{10} x_{m+j},$$
where the second follows from $\Sigma^A_{m_1+j,m+j} = 9/20$, $\Sigma^A_{m_1+j,r} = 1/2$ for all $r \in [n] \backslash \{m_1+j, m+j\}$, and $\Sigma^B_{m_1+j,r} = 0$ for all $r \in [n] \backslash \{m_1+j\}$. Hence, by Lemma 3.4, we have
$$\begin{aligned}\tau_{m_1+j} &= \left(\frac{\Delta_{m_1+j,1} - \Delta_{m_1+j,2}}{\Delta_{m_1+j,1}^2} \pm O\left(\frac{\beta^2}{\Delta_{m_1+j,1}}\right)\right) \cdot b_2 \\ &= \left(\frac{1}{\Delta_{m_1+j,1}} - \frac{1}{10} \cdot \frac{x_{m+j}}{\Delta_{m_1+j,1}^2} \pm O\left(\frac{\beta^2}{\Delta_{m_1+j,1}}\right)\right) \cdot b_2\end{aligned}$$
Given that $\mathcal{D}_{m_1+j}$ is uniform over
$$\left[\left(\frac{1}{\Delta_{m_1+j,1}} - \frac{1}{10} \cdot \frac{\beta\delta}{\Delta_{m_1+j,1}^2}\right) \cdot b_2, \frac{1}{\Delta_{m_1+j,1}} \cdot b_2\right]$$
with mass $\beta\delta$, we have
$$x_{m_1+j} = \mathsf{T}_{[0,\beta\delta]}(\beta\delta - x_{m+j} \pm O(\kappa\beta\delta)). \tag{14}$$

It remains to analyse the jumping probability of player $m+j$, and we have

$$\begin{aligned}\Delta_{m+j,2} &= \sum_{r\in[n]\backslash\{m+j\}} \left(1 - 2\Sigma^A_{m+j,r} - \Sigma^B_{m+j,r}\right) x_r \\ &= -\frac{1}{10}x_{m_1+j} - \frac{1}{5}x_{j(1)} \\ &= -\frac{1}{10}\left(\beta\delta - x_{m+j} + 2x_{j(1)}\right) \pm O(\kappa\beta\delta).\end{aligned}$$

Here the second step follows from $\Sigma^A_{m+j,m_1+j} = 11/20$, $\Sigma^A_{m+j,r} = 1/2$ for $r \in [n]\backslash\{m_1+j, m+j\}$, $\Sigma^B_{m+j,j(1)} = 1/5$ and $\Sigma^B_{m+j,r} = 0$ for any $r \in [n]\backslash\{m+j, j(1)\}$. The last step follows from Eq. (14).

Now, by Lemma 3.4, one has

$$\begin{aligned}&\tau_{m+j} \\ &= \left(\frac{\Delta_{m+j,1} - \Delta_{m+j,2}}{\Delta^2_{m+j,1}} \pm O\left(\frac{\beta^2}{\Delta_{m+j,1}}\right)\right) \cdot b_2 \\ &= \left(\frac{1}{\Delta_{m+j,1}} + \frac{1}{10} \cdot \frac{\beta\delta - x_{m+j} + 2x_{j(1)} \pm O(\kappa\beta\delta)}{\Delta^2_{m+j,1}} \pm O\left(\frac{\beta^2}{\Delta_{m+j,1}}\right)\right) \cdot b_2\end{aligned}$$

Given that $\mathcal{D}_{m+j}$ is uniform over

$$\left[\left(\frac{1}{\Delta_{m+j,1}} + \frac{1}{10}\cdot\frac{\beta\delta}{\Delta^2_{m+j,1}}\right)\cdot b_2, \left(\frac{1}{\Delta_{m+j,1}} + \frac{1}{5}\cdot\frac{\beta\delta}{\Delta^2_{m+j,1}}\right)\cdot b_2\right]$$

with mass $\beta\delta$, we conclude that

$$x_{m+j} = \mathsf{T}_{[0,\beta\delta]}(x_{j(1)} \pm O(\kappa\beta\delta)).$$

Plugging into Eq. (14), we obtain

$$x_{m_1+j} = \mathsf{T}_{[0,\beta\delta]}(\beta\delta - x_{j(1)} \pm O(\kappa\beta\delta)).$$

This completes the proof of the second claim. □

4 PTAS

We present a PTAS for computing an ϵ-approximate BNE in an FPA under the uniform tie-breaking rule.

Theorem 1.2 (PTAS under uniform tie-breaking). *For any $\epsilon > 0$, $n, m \geq 2$, there is an algorithm that finds an ϵ-approximate Bayesian Nash equilibrium using $O(n^4 \cdot g(1/\epsilon))$ time under the uniform tie-breaking rule.*

Our approach proceeds in the following four steps. Given an FPA $(\mathcal{N}, \mathcal{B}, \mathcal{D}, \Gamma)$ where Γ is the uniform tie-breaking rule, we first round the bidding space $\mathcal{B}$ and reduce its size to $O(1/\epsilon)$. We then prune the valuation distribution $\mathcal{D}$ and work on a weak notion of (ϵ, δ)-approximate BNE that relaxes the no-overbidding requirement. In the third step, we argue the existence of an (ϵ, δ)-approximate BNE profile over a discretized space and in the final step, we develop a suitable searching algorithm for (ϵ, δ)-approximate BNE in the discretized space. Missing proof can be found at the full version of this paper.

Step 1: Rounding bids. Given a first-price auction with bidding space $\mathcal{B} = \{b_0, b_1, \ldots, b_m\}$ with $0 = b_0 < b_1 < \cdots < b_m \leq 1$, we define $\mathcal{B}_\epsilon = \{b_{0,\max}, b_{1,\max}, \ldots, b_{10\epsilon^{-1},\max}\}$ as follows. First we take $b_{0,\max} = b_0 = 0$. Then for each $t \in [10\epsilon^{-1}]$, let $b_{t,\max}$ be the maximum bid in

$$\mathcal{B}_t := \mathcal{B} \cap \left(\frac{(t-1)\epsilon}{10}, \frac{t\epsilon}{10}\right]$$

if $\mathcal{B}_t$ is not empty; and set $b_{t,\max} = \mathsf{nil}$ (meaning that we don't add an element to $\mathcal{B}_\epsilon$) if $\mathcal{B}_t$ is empty. We prove that it suffices to find an $(\epsilon/2)$-approximate BNE over bidding space $\mathcal{B}_\epsilon$.

Lemma 4.1. *Given any first-price auction $(\mathcal{N}, \mathcal{B}, \mathcal{D}, \Gamma)$, let B_ϵ be the rounded bidding space defined above. Then any $(\epsilon/2)$-approximate BNE of $(\mathcal{N}, \mathcal{B}_\epsilon, \mathcal{D}, \Gamma)$ is also an ϵ-approximate BNE of $(\mathcal{N}, \mathcal{B}, \mathcal{D}, \Gamma)$.*

Step 2: Rounding distribution. Given a first-price auction $(\mathcal{N}, \mathcal{B}_\epsilon, \mathcal{D}, \Gamma)$, we would like to round the value distribution $\mathcal{D}_i$ such that it is supported over discrete values, and we truncate off the probability mass if it is too small. Formally, letting $\delta \in (0, \frac{\epsilon}{20})$ be a parameter to be specified later[4], we define $\mathcal{D}_i^{\epsilon,\delta}$ for each $i \in [n]$ as follows:

$$\begin{aligned}p_i^{\epsilon,\delta}(t) &= \Pr_{v_i \sim \mathcal{D}_i^{\epsilon,\delta}}\left[v_i = \frac{t\epsilon}{10}\right] \\ &= \max\left\{0, \int_{\frac{t\epsilon}{10}}^{\frac{(t+1)\epsilon}{10}} \Pr_{v_i\sim\mathcal{D}_i}[v_i = v]\,\mathrm{d}v - \delta\right\}\end{aligned}$$

for each $t \in [10\epsilon^{-1} - 1]$, and define

$$p_i^{\epsilon,\delta}(0) = \Pr_{v_i\sim\mathcal{D}_i^{\epsilon,\delta}}\left[v_i = 0\right] = 1 - \sum_{t=1}^{10\epsilon^{-1}-1} \Pr_{v_i\sim\mathcal{D}_i^{\epsilon,\delta}}\left[v_i = \frac{t\epsilon}{10}\right].$$

That is, the valuation distribution is rounded (down) to discrete values $V_\epsilon := \{0, \epsilon/10, \ldots, 1 - \epsilon/10\}$ and truncated at δ; the extra probability mass is put on 0. From now on, we would consider both continuous and discrete valuation distribution of bidders.

Let $\mathcal{D}^{\epsilon,\delta} = \mathcal{D}_1^{\epsilon,\delta} \times \cdots \times \mathcal{D}_n^{\epsilon,\delta}$. Next we define the notion of (ϵ,δ)-approximate BNE — a weaker notation of equilibrium with relaxed constraint on overbidding, and show that it suffices to find an (ϵ,δ)-approximate BNE of the rounded FPA $(\mathcal{N}, \mathcal{B}_\epsilon, \mathcal{D}^{\epsilon,\delta}, \Gamma)$.

Definition 4.2 ((ϵ,δ)-approximate Bayesian Nash equilibrium)**.** *Let $n, m \geq 2$. Given a Bayesian FPA $(\mathcal{N}, \mathcal{B}_\epsilon, \mathcal{D}^{\epsilon,\delta}, \Gamma)$, a strategy profile $s = (s_1, \ldots, s_n)$ is said to be an (ϵ,δ)-approximate Bayesian Nash equilibrium ((ϵ,δ)-approximate BNE), if for any player $i \in [n]$, its strategy s_i is monotone and at most ϵ-worse than the best response:*

$$\begin{aligned}&\mathop{\mathbb{E}}_{v\sim\mathcal{D}}\left[u_i(v_i; s_i(v_i), s_{-i}(v_{-i}))\right] \\ \geq\ &\mathop{\mathbb{E}}_{v\sim\mathcal{D}}\left[u_i(v_i; \mathrm{bs}(v_i, s_{-i}), s_{-i}(v_{-i}))\right] - \epsilon,\end{aligned}$$

Moreover, letting $v_{i,\max} = \max_{v_i\in V_\epsilon, p_i(v_{i,\max})>0} v_i$, the player i never bids higher than $v_{i,\max}$ and its total overbidding probability is at most δ, i.e.,

$$\Pr_{v_i\sim\mathcal{D}_i^{\epsilon,\delta}}\left[s_i(v_i) > v_{i,\max}\right] = 0 \quad \textit{and} \quad \Pr_{v_i\sim\mathcal{D}_i^{\epsilon,\delta}}\left[s_i(v_i) > v_i\right] < \delta.$$

We prove that it suffices to find an (ϵ,δ)-approximate BNE in the rounded FPA $(\mathcal{N}, \mathcal{B}_\epsilon, \mathcal{D}^{\epsilon,\delta}, \Gamma)$. In the proof we will perform the operation of converting a bidding strategy s_i over bidding space $\mathcal{B}_\epsilon$ and value distribution $\mathcal{D}_i$ to a (unique) monotone bidding strategy

[4]Looking head, the parameter δ shall be much smaller than ϵ

s_i' over the same bidding space $\mathcal{B}_\epsilon$ but a different value distribution $\mathcal{D}_i'$ such that their induced distributions over $\mathcal{B}_\epsilon$ are the same:

Definition 4.3 (Monotone analogue). *Let $s = (s_1, \ldots, s_n)$ be a strategy profile over bidding space $\mathcal{B}_\epsilon$ and value distribution $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$, and let $\mathcal{D}' = \mathcal{D}_1' \times \cdots \times \mathcal{D}_n'$ be another value distribution. The* monotone analogue *of s with respect to $\mathcal{D}'$ (denoted as s^{m}) is defined as the unique monotone strategy profile $s^{\mathrm{m}} = (s_1^{\mathrm{m}}, \ldots, s_n^{\mathrm{m}})$ such that*

$$\Pr_{v_i \sim \mathcal{D}_i'}\left[s_i^{\mathrm{m}}(v_i) = b\right] = \Pr_{v_i \sim \mathcal{D}_i}\left[s_i(v_i) = b\right], \quad \text{for all } i \in [n] \text{ and } b \in \mathcal{B}_\epsilon.$$

Lemma 4.4. *Given any FPA $(\mathcal{N}, \mathcal{B}_\epsilon, \mathcal{D}, \Gamma)$, let $\mathcal{D}^{\epsilon,\delta}$ be the rounded distribution defined above. Then the monotone analogue of any (ϵ, δ)-approximate BNE in $(\mathcal{N}, \mathcal{B}_\epsilon, \mathcal{D}^{\epsilon,\delta}, \Gamma)$ is an $(2\epsilon + 22\delta/\epsilon)$-approximate BNE in $(\mathcal{N}, \mathcal{B}_\epsilon, \mathcal{D}, \Gamma)$.*

Step 3: Existence of discretized (ϵ, δ)-approximate BNE. Given a first-price auction $(\mathcal{N}, \mathcal{B}, \mathcal{D}, \Gamma)$, we prove the existence of a (suitably) discretized (ϵ, δ)-approximate BNE. We describe this step using a generic FPA $(\mathcal{N}, \mathcal{B}, \mathcal{D}, \Gamma)$ but will apply it on the rounded FPA $(\mathcal{N}, \mathcal{B}_\epsilon, \mathcal{D}^{\epsilon,\delta}, \Gamma)$ later. For any $j \in [0 : m]$, $i \in [n]$, let $p_{i,j} = \Pr[s_i(v_i) = b_j]$ be the probability of player i bidding b_j in a strategy profile s.

Lemma 4.5 (Discretization). *Let $m, n \geq 2$, given any first-price auction $(\mathcal{N}, \mathcal{B}, \mathcal{D}, \Gamma)$ and $\mathcal{B} = \{b_0, b_1, \ldots, b_m\}$ with $0 = b_0 < b_1 < \cdots < b_m \leq 1$, there exists an (ϵ, δ)-approximate BNE strategy profile s such that $p_{i,j}$ is a integer multiple of*

$$\ell(m) \cdot \frac{1}{\epsilon^6 \delta}$$

for any $i \in [n]$ and $j \in [0 : m]$, where $\ell(m)$ is some exponential function of m.

We make use of the following result from [18]. We note this the major part that we need a uniform tie-breaking rule.

Theorem 4.6 (Theorem 3 of [18]). *Let $p_i \in \Delta_{m+1}$ for $i \in [n]$, and let $\{X_i \in \mathbb{R}^{m+1}\}_{i \in [n]}$ be a set of independent $(m+1)$-dimensional random unit vectors such that, for all $i \in [n]$, $j \in [0 : m]$, $\Pr[X_i = 1_j] = p_{i,j}$. Let $z > 0$ be an integer. Then there exists another set of probability vectors $\{\widehat{p}_i \in \Delta_{m+1}\}_{i \in [n]}$ such that the following conditions hold:*

- *$|\widehat{p}_{i,j} - p_{i,j}| = O(1/z)$, for all $i \in [n]$ and $j \in [0 : m]$;*
- *$\widehat{p}_{i,j}$ is an integer multiple of $\frac{1}{2^m} \cdot \frac{1}{z}$ for all $i \in [n]$ and $j \in [0 : m]$;*
- *If $p_{i,j} = 0$ then $\widehat{p}_{i,j} = 0$;*
- *Let $\{\widehat{X}_i \in \mathbb{R}^{m+1}\}_{i \in [n]}$ be a set of independent $(m+1)$-dimensional random unit vectors such that $\Pr[X_i = 1_j] = \widehat{p}_{i,j}$ for all $i \in [n]$, $j \in [0 : m]$.Then*

$$\left\| \sum_{i \in [n]} X_i - \sum_{i \in [n]} \widehat{X}_i \right\|_{\mathsf{TV}} = O\left(h(m) \cdot \frac{\log z}{z^{1/5}}\right). \tag{15}$$

Moreover, for all $i' \in [n]$, we have

$$\left\| \sum_{i \in [n] \setminus \{i'\}} X_i - \sum_{i \in [n] \setminus \{i'\}} \widehat{X}_i \right\|_{\mathsf{TV}} = O\left(h(m) \cdot \frac{\log z}{z^{1/5}}\right). \tag{16}$$

where $h(m)$ is some exponential function

Proof of Lemma 4.5. Given a BNE strategy profile s of FPA $(\mathcal{N}, \mathcal{B}, \mathcal{D}, \Gamma)$, we take

$$z = \Omega\Big(\max\left\{2h(m)^6 \epsilon^{-6}, 2m^2 \delta^{-1}\right\}\Big) \quad \text{and} \quad \gamma = h(m) \cdot \frac{\log z}{z^{1/5}}.$$

Let $\{p_i \in \Delta^{m+1}\}_{i \in [n]}$ be the probability vectors that correspond to s, i.e., $p_{i,j} = \Pr[s_i(v_i) = b_j]$. Using Theorem 4.6, let $\{\widehat{p}_i \in \Delta^{m+1}\}_{i \in [n]}$ be the set of discretized probability vectors, and let $\widehat{s}$ be the unique monotone strategy determined by $\widehat{p}$ (with respect to the same value distribution $\mathcal{D}$). We prove that $\widehat{s}$ forms an (ϵ, δ)-approximate BNE of the auction. We need to verify that for every player $i \in [n]$, (1) $\widehat{s}_i$ is at most ϵ-worse than the best response; and (2) the overbidding probability is small.

By Eq. (16), for any $j \in [0 : m]$, one has

$$\mathop{\mathbb{E}}_{v_{-i} \sim \mathcal{D}_{-i}}\left[\Gamma_i(b_j, s_{-i}(v_{-i}))\right] = \mathop{\mathbb{E}}_{v_{-i} \sim \mathcal{D}_{-i}}\left[\Gamma_i(b_j, \widehat{s}_{-i}(v_{-i}))\right] \pm \gamma. \tag{17}$$

since the allocation probability (under the uniform tie-breaking) is determined by the histogram of other players' bidding histogram, which shifts by at most γ between s_{-i} and $\widehat{s}_{-i}$ in total variance distance.

For the utility, we have

$$\begin{aligned}
&\mathop{\mathbb{E}}_{v \sim \mathcal{D}}\left[u_i(v_i; \widehat{s}_i(v_i), \widehat{s}_{-i}(v_{-i}))\right] \\
&= \mathop{\mathbb{E}}_{v \sim \mathcal{D}}\left[(v_i - \widehat{s}_i(v_i)) \cdot \Gamma_i(\widehat{s}(v_i), \widehat{s}_{-i}(v_{-i}))\right] \\
&= \mathop{\mathbb{E}}_{v \sim \mathcal{D}}\left[(v_i - \widehat{s}_i(v_i)) \cdot \Gamma_i(\widehat{s}_i(v_i), s_{-i}(v_i))\right] \pm \gamma \\
&= \mathop{\mathbb{E}}_{v \sim \mathcal{D}}\left[(v_i - s_i(v_i)) \cdot \Gamma_i(s_i(v_i), s_{-i}(v_i))\right] \pm O\left(m^2 \cdot \frac{1}{z} + \gamma\right) \\
&= \mathop{\mathbb{E}}_{v \sim \mathcal{D}}\left[u_i(v_i; s(v_i), s_{-i}(v_{-i}))\right] \pm \frac{\epsilon}{2}.
\end{aligned} \tag{18}$$

The second step follows from Eq. (17), the third step holds since the TV distance between $(v_i, s_i(v_i))$ and $(v_i, \widehat{s}_i(v_i))$ is at most $O(m^2 \cdot \frac{1}{z})$. The last step follows from the choice of z.

At the same time, we have

$$\begin{aligned}
&\mathop{\mathbb{E}}_{v \sim \mathcal{D}}\left[u_i(v_i; \mathrm{bs}(v_i, \widehat{s}_{-i}), \widehat{s}_{-i}(v_{-i}))\right] \\
&\leq \mathop{\mathbb{E}}_{v \sim \mathcal{D}}\left[u_i(v_i, \mathrm{bs}(v_i, \widehat{s}_{-i}), s_{-i}(v_{-i}))\right] + \gamma \\
&\leq \mathop{\mathbb{E}}_{v \sim \mathcal{D}}\left[u_i(v_i, \mathrm{bs}(v_i, s_{-i}), s_{-i}(v_{-i}))\right] + \gamma \\
&\leq \mathop{\mathbb{E}}_{v \sim \mathcal{D}}\left[u_i(v_i; \mathrm{bs}(v_i, s_{-i}), s_{-i}(v_{-i}))\right] + \frac{\epsilon}{2}.
\end{aligned} \tag{19}$$

Here the first step follows from the bidding histogram shifts by at most γ between s_{-i} and $\widehat{s}_{-i}$ in total variance distance. The last step follows from the choice of parameters.

Combining Eq. (18) and Eq. (19), we have that $\widehat{s}_i$ is at most ϵ-worse than the best response.

To bound the probability of overbidding, let $j(i) \in [0 : m]$ be the maximum bid that receives non-zero probability under $\widehat{s}_i$, then it is easy to verify that $b_{j(i)} \leq v_{i,\max}$: otherwise $b_{j(i)} > v_{i,\max}$ and $p_{i,j(i)} > 0$ would contradict with the no overbidding assumption of BNE on s.

Finally we have

$$\Pr_{v_i\sim\mathcal{D}_i}\left[\widehat{s}_i(v_i) > v_i\right] \leq \Pr_{v_i\sim\mathcal{D}_i}\left[s_i(v_i) > v_i\right] + O\left(m^2\cdot\frac{1}{z}\right) = 0 + O\left(m^2\cdot\frac{1}{z}\right) \leq \delta$$

since the total variation distance between $(v_i, s_i(v_i))$ and $(v_i, \widehat{s}_i(v_i))$ is at most $O(m^2 \cdot \frac{1}{z})$. □

Step 4: Searching for an (ϵ,δ)-approximate BNE. Finally we provide a simple searching algorithm for (ϵ,δ)-approximate BNE over the discretized space. The key observation comes from the allocation rule of a first-price auction, i.e., only players with the highest bid could win the item. We say a strategy profile s lies in the grid S_ω for some $\omega \in (0,1)$ if $p_{i,j} = \Pr_{v_i\sim\mathcal{D}_i}[s_i(v_i) = b_j]$ is a multiple of ω for every $i \in [n]$ and $j \in [0:m]$.

Lemma 4.7. *Given a first-price auction $(\mathcal{N},\mathcal{B},\mathcal{D},\Gamma)$, and suppose there exists at least one (ϵ,δ)-approximate BNE over the grid S_ω, then there is an algorithm that runs in $n^4m\cdot 2^{\widetilde{O}(m/\epsilon\omega)}$ time and returns an $(2\epsilon,\delta)$-approximate BNE under the uniform tie-breaking rule.*

Proof. Let $R = 100(m+1)/\epsilon\omega$. The discretized strategy profiles $\mathsf{E}_\omega \subseteq [\Delta_{m+1}]^R$ are defined as follows. A strategy profile $(p_1,\dots,p_R) \in \mathsf{E}_\omega$ is parameterized by a bid level $j^* \in [0:m-1]$ and $k_0, k_1, \dots, k_{j^*} \in [0:10\epsilon^{-1}]$ such that $p_{i,j}$ is a multiple of ω for all $i \in [n], j \in [0:m]$, and

$$\sum_{r=1}^{R} p_{r,m-j} = k_j,\ \forall j \in [0:j^*]$$
$$p_{r,m-j} = 0,\ \forall r \in [R], j \in [j^*+1:m-1].$$

The size of E_ω satisfies

$$|\mathsf{E}_\omega| \leq (m+1)\cdot(10/\epsilon\omega + 1/\omega)^{m+1}\cdot\binom{R}{10/\epsilon\omega+1/\omega} \leq 2^{\widetilde{O}(m/\epsilon\omega)}.$$

For any strategy profile $(p_1,\dots,p_R) \in \mathsf{E}_\omega$, one can augment with $(n-R)$ default strategies $(1,0,\dots,0)$ (that is, bidding $b_0 = 0$ with probability 1). Slightly abuse of notation, we also use $\mathsf{E}_\omega \subseteq [\Delta_{m+1}]^n$ to denote the augmented strategy profiles. We shall prove

- There is an $(2\epsilon,\delta)$-approximate BNE in E_ω (up to a matching with players), and
- One can identify the matching efficiently.

Existence of $(2\epsilon,\delta)$-approximate BNE. Recall that there exists an (ϵ,δ)-approximate BNE strategy s over the grid S_ω. Let j^* be the first index over $[0:m]$ such that $\sum_{i\in[n]} p_{i,m-j} \leq 10/\epsilon$ $(\forall j < j^*)$ and $\sum_{i\in[n]} p_{i,m-j^*} \geq 10/\epsilon$. If $j^* = m$, then we have $s \in \mathsf{E}_\omega$ (up to a matching between players). If $j^* \leq m-1$, then let $n^* \in [n]$ be the smallest player such that $\sum_{i\in[n^*]} p_{n,m-j^*} \geq 10/\epsilon$ (w.l.o.g. we assume it takes the equality). Truncate the strategy profile to s' such that

$$p'_{i,j} = \Pr_{v_i\sim\mathcal{D}_i}[s'_i(v_i) = b_j] = \begin{cases} p_{i,j} & j > m-j^* \vee (j = m-j^* \wedge i \leq n^*) \\ 0 & j \in [1:m-j^*-1] \vee (j = m-j^* \wedge i > n^*) \\ 1 - \sum_{j'=1}^{m} p'_{i,j'} & j = 0. \end{cases}$$

It is clear the new strategy $s' \in \mathsf{E}_\omega$ (up to a matching between players), and for each player $i \in [n]$, the new strategy s'_i is monotone and the probability of overbidding is no more than δ (because we only move bidding probability to $b_0 = 0$). It suffices to prove it is at most 2ϵ-worse than the best response. The key observation is that the allocation probability of bidding b_j with $j > m-j^*$ remains the same, i.e., for any $j \in [m-j^*+1:m]$

$$\mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}[\Gamma_i(b_j, s'_{-i}(v_{-i}))] = \mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}[\Gamma_i(b_j, s_{-i}(v_{-i}))]$$

and moreover, the allocation probability of bidding no more than b_{m-j^*} is small

$$\mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}[\Gamma_i(b_j; s'_{-i}(v_{-i}))] \leq \epsilon/4 \quad \forall j \in [0:m-j^*].$$

This holds since with probability at least $1-\exp(-5/3\epsilon)$, there are at least $5/\epsilon$ players bid no less than b_{m-j^*} by Chernoff bound.

Therefore, at any value point v_i, if $s'_i(v_i) > b_{m-j^*}$, then $s_i(v_i) = s'_i(v_i)$ and

$$\mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}[u_i(v_i; s'_i(v_i); s'_{-i}(v_{-i}))] = \mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}[u_i(v_i; s_i(v_i); s_{-i}(v_{-i}))] \tag{20}$$

If $s'_i(v_i) \leq b_{m-j^*}$, then $s_i(v_i) \leq b_{m-j^*}$, and

$$\mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}[u_i(v_i; s'_i(v_i); s'_{-i}(v_{-i}))] \geq -\epsilon/4$$
$$\text{and}\quad \mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}[u_i(v_i; s_i(v_i); s_{-i}(v_{-i}))] \leq \epsilon/4. \tag{21}$$

Combining Eq. (20) and Eq. (21), for any $v_i \in [0,1]$, one has

$$\mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}[u_i(v_i; s'_i(v_i); s'_{-i}(v_{-i})) - u_i(v_i; s_i(v_i); s_{-i}(v_{-i}))] \geq -\epsilon/2.$$

Similarly, one can prove

$$\mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}[u_i(v_i; \mathrm{bs}(v_i, s'_{-i}), s'_{-i}(v_{-i}))] \leq \mathop{\mathbb{E}}_{v_{-i}\sim\mathcal{D}_{-i}}[u_i(v_i; \mathrm{bs}(v_i, s_{-i}), s_{-i}(v_{-i}))] + \epsilon/2.$$

Combining the above two inequalities, we have proved s'_i is at most 2ϵ-worse than the best response.

Find a matching. Given a strategy profile $(p_1,\dots,p_n) \in \mathsf{E}_\omega$, we show how to define a bipartite matching problem, such that the $(2\epsilon,\delta)$-approximate BNE are one-to-one correspondence to the perfect bipartite matching of the graph. The bipartite matching problem is defined between players $[n]$ and the strategy $\{p_i\}_{i\in[n]}$. We draw an edge between player i_1 and the strategy p_{i_2}, if p_{i_2} is at most 2ϵ-worse than the best response (note the histogram of other players' bidding profile is determined) and the overbidding probability is small at most δ. We note the best response can be computed in $O(mn^2)$ time and one can estimate the utility of a bid in $O(n^2)$ time using dynamic programming (similar as [24]). Hence, it takes $n^2\cdot O(mn^2)$ time to construct the bipartite graph, and a perfect matching can be found in time $O(n^3)$. □

ACKNOWLEDGEMENT

X.C. and B.P. would like to thank Aviad Rubinstein and anonymous STOC reviewers for helpful suggestions on the paper. The research of X.C. and B.P. is supported by NSF grants CCF-1703925, IIS-1838154, CCF-2106429, CCF-2107187, CCF-1763970, CCF-2212233, COLL2134095, COLL2212745.

REFERENCES

[1] Tim Abbott, Daniel Kane, and Paul Valiant. 2005. On the complexity of two-player win-lose games. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*. IEEE, 113–122.

[2] Susan Athey. 2001. Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica* 69, 4 (2001), 861–889.

[3] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. 2017. First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica* 85, 1 (2017), 107–143.

[4] Yang Cai and Christos Papadimitriou. 2014. Simultaneous bayesian auctions and computational complexity. In *Proceedings of the fifteenth ACM conference on Economics and computation*. 895–910.

[5] Bhaskar Ray Chaudhury, Jugal Garg, Peter McGlaughlin, and Ruta Mehta. 2021. Competitive allocation of a mixed manna. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1405–1424.

[6] Shuchi Chawla and Jason D Hartline. 2013. Auctions with unique equilibria. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. 181–196.

[7] Thomas Chen, Xi Chen, Binghui Peng, and Mihalis Yannakakis. 2022. Computational Hardness of the Hylland-Zeckhauser Scheme. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2253–2268.

[8] Xi Chen, Decheng Dai, Ye Du, and Shang-Hua Teng. 2009. Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, 273–282.

[9] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. 2009. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM (JACM)* 56, 3 (2009), 1–57.

[10] Xi Chen, David Durfee, and Anthi Orfanou. 2015. On the complexity of Nash equilibria in anonymous games. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*. 381–390.

[11] Xi Chen, Dimitris Paparas, and Mihalis Yannakakis. 2013. The complexity of non-monotone markets. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. 181–190.

[12] Xi Chen and Shang-Hua Teng. 2009. Spending is not easier than trading: on the computational equivalence of Fisher and Arrow-Debreu equilibria. In *International Symposium on Algorithms and Computation*. Springer, 647–656.

[13] Xi Chen, Shang-Hua Teng, and Paul Valiant. 2007. The approximation complexity of win-lose games. In *SODA*, Vol. 7. 159–168.

[14] Yeshwanth Cherapanamjeri, Constantinos Daskalakis, Andrew Ilyas, and Manolis Zampetakis. 2022. Estimation of Standard Auction Models. In *Proceedings of the 23rd ACM Conference on Economics and Computation*.

[15] Vincent Conitzer, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Nicolas E Stier-Moses, Eric Sodomka, and Christopher A Wilkens. 2022. Pacing Equilibrium in First Price Auction Markets. *Management Science* (2022).

[16] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. 2009. The complexity of computing a Nash equilibrium. *SIAM J. Comput.* 39, 1 (2009), 195–259.

[17] Constantinos Daskalakis, Noah Golowich, and Kaiqing Zhang. 2022. The complexity of markov equilibrium in stochastic games. *arXiv preprint arXiv:2204.03991* (2022).

[18] Constantinos Daskalakis and Christos H Papadimitriou. 2015. Approximate Nash equilibria in anonymous games. *Journal of Economic Theory* 156 (2015), 207–245.

[19] Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. 2021. The complexity of constrained min-max optimization. In *Proceedings of the 53th Annual ACM SIGACT Symposium on Theory of Computing*.

[20] Argyrios Deligkas, John Fearnley, Alexandros Hollender, and Themistoklis Melissourgos. 2022. Pure-circuit: Strong inapproximability for PPAD. In *2022 IEEE 63th Annual Symposium on Foundations of Computer Science (FOCS)*.

[21] Argyrios Deligkas, John Fearnley, Alexandros Hollender, and Themistoklis Melissourgos. 2022. Pure-Circuit:Strong Inapproximability for PPAD. In *53th Annual IEEE Symposium on Foundations of Computer Science (FOCS'22)*. IEEE.

[22] Yuan Deng, Jieming Mao, Vahab Mirrokni, Hanrui Zhang, and Song Zuo. 2022. Efficiency of the first-price auction in the autobidding world. *arXiv preprint arXiv:2208.10650* (2022).

[23] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. 2013. Simultaneous auctions are (almost) efficient. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. 201–210.

[24] Aris Filos-Ratsikas, Yiannis Giannakopoulos, Alexandros Hollender, Philip Lazos, and Diogo Poças. 2021. On the complexity of equilibrium computation in first-price auctions. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 454–476.

[25] Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. 2000. Optimal nonparametric estimation of first-price auctions. *Econometrica* 68, 3 (2000), 525–574.

[26] Jason Hartline, Darrell Hoy, and Sam Taggart. 2014. Price of anarchy for auction revenue. In *Proceedings of the fifteenth ACM conference on Economics and computation*. 693–710.

[27] Darrell Hoy, Samuel Taggart, and Zihe Wang. 2018. A tighter welfare guarantee for first-price auctions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. 132–137.

[28] Yaonan Jin and Pinyan Lu. 2022. First Price Auction is $1 - 1/e^2$ Efficient. In *2022 IEEE 63th Annual Symposium on Foundations of Computer Science (FOCS)*.

[29] Yaonan Jin and Pinyan Lu. 2023. The Price of Stability for First Price Auction. In *SODA*.

[30] Yujia Jin, Vidya Muthukumar, and Aaron Sidford. 2022. The Complexity of Infinite-Horizon General-Sum Stochastic Games. *arXiv preprint arXiv:2204.04186* (2022).

[31] Todd R Kaplan and Shmuel Zamir. 2012. Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory* 50, 2 (2012), 269–302.

[32] Bernard Lebrun. 1996. Existence of an equilibrium in first price auctions. *Economic Theory* 7, 3 (1996), 421–443.

[33] Bernard Lebrun. 1999. First price auctions in the asymmetric N bidder case. *International Economic Review* 40, 1 (1999), 125–142.

[34] Bernard Lebrun. 2006. Uniqueness of the equilibrium in first-price auctions. *Games and Economic Behavior* 55, 1 (2006), 131–151.

[35] Alessandro Lizzeri and Nicola Persico. 2000. Uniqueness and existence of equilibrium in auctions with a reserve price. *Games and Economic Behavior* 30, 1 (2000), 83–114.

[36] Robert C Marshall, Michael J Meurer, Jean-Francois Richard, and Walter Stromquist. 1994. Numerical analysis of asymmetric first price auctions. *Games and Economic behavior* 7, 2 (1994), 193–220.

[37] Eric Maskin and John Riley. 2000. Equilibrium in sealed high bid auctions. *The Review of Economic Studies* 67, 3 (2000), 439–454.

[38] Eric Maskin and John Riley. 2003. Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior* 45, 2 (2003), 395–409.

[39] Ruta Mehta. 2014. Constant rank bimatrix games are PPAD-hard. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*. 545–554.

[40] Abraham Othman, Christos Papadimitriou, and Aviad Rubinstein. 2016. The complexity of fairness through equilibrium. *ACM Transactions on Economics and Computation (TEAC)* 4, 4 (2016), 1–19.

[41] Christos Papadimitriou and Binghui Peng. 2021. Public goods games in directed networks. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 745–762.

[42] Christos H Papadimitriou. 1994. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and system Sciences* 48, 3 (1994), 498–532.

[43] Michael Plum. 1992. Characterization and computation of Nash-equilibria for auctions with incomplete information. *International Journal of Game Theory* 20, 4 (1992), 393–418.

[44] Philip J Reny and Shmuel Zamir. 2004. On the existence of pure strategy monotone equilibria in asymmetric first-price auctions. *Econometrica* 72, 4 (2004), 1105–1125.

[45] John G Riley and William F Samuelson. 1981. Optimal auctions. *The American Economic Review* 71, 3 (1981), 381–392.

[46] Tim Roughgarden. 2010. Algorithmic game theory. *Commun. ACM* 53, 7 (2010), 78–86.

[47] Aviad Rubinstein. 2015. Inapproximability of Nash equilibrium. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*. 409–418.

[48] Aviad Rubinstein. 2016. Settling the complexity of computing approximate two-player Nash equilibria. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 258–265.

[49] Aviad Rubinstein. 2018. Inapproximability of Nash equilibrium. *SIAM J. Comput.* 47, 3 (2018), 917–959.

[50] Aviad Rubinstein. 2019. *Hardness of Approximation Between P and NP*. Morgan & Claypool.

[51] Vasileios Syrgkanis. 2014. *Efficiency of mechanisms in complex markets*. Cornell University.

[52] Vasilis Syrgkanis and Eva Tardos. 2013. Composable and efficient mechanisms. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. 211–220.

[53] Vijay V Vazirani and Mihalis Yannakakis. 2011. Market equilibrium under separable, piecewise-linear, concave utilities. *Journal of the ACM (JACM)* 58, 3 (2011), 1–25.

[54] William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance* 16, 1 (1961), 8–37.

[55] Zihe Wang, Weiran Shen, and Song Zuo. 2020. Bayesian nash equilibrium in first-price auction with discrete value distributions. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. 1458–1466.