

On the Computation of Equilibria in Discrete First-Price Auctions*

Aris Filos-Ratsikas

University of Edinburgh, United Kingdom
aris.filos-ratsikas@ed.ac.uk

Yiannis Giannakopoulos

University of Glasgow, United Kingdom
yiannis.giannakopoulos@glasgow.ac.uk

Alexandros Hollender

University of Oxford, United Kingdom
alexandros.hollender@cs.ox.ac.uk

Charalampos Kokkalis

University of Edinburgh, United Kingdom
charalampos.kokkalis@ed.ac.uk

Abstract

We study the computational complexity of computing Bayes-Nash equilibria in first-price auctions with discrete value distributions and discrete bidding space, under general subjective beliefs. It is known that such auctions do not always have pure equilibria. In this paper, we prove that the problem of deciding their existence is NP-complete, even for approximate equilibria. On the other hand, it can be shown that *mixed* equilibria are guaranteed to exist; however, their computational complexity has not been studied before. We establish the PPAD-completeness of computing a mixed equilibrium and we complement this by an efficient algorithm for finding symmetric approximate equilibria in the special case of iid priors. En route to these results, we develop a computational equivalence framework between continuous and discrete first-price auctions, which can be of independent interest, and which allows us to transfer existing positive and negative results from one setting to the other. Finally, we show that correlated equilibria of the auction can be computed in polynomial time.

*Aris Filos-Ratsikas was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/Y003624/1. Charalampos Kokkalis was supported by an EPSRC DTA Scholarship (Reference EP/W524384/1). Some preliminary work on this project was performed during a visit of Alexandros Hollender to the University of Edinburgh under the SICSA Distinguished Visiting Fellowship scheme.

A preliminary version of this paper appeared in EC'24.

Contents

1	Introduction	3
1.1	Our Results and Discussion	4
1.2	Further Related Work	6
2	Preliminaries	6
2.1	The Auction Game	7
2.2	Efficient Computation of Bidder Utilities	9
3	Pure Equilibria	10
4	Mixed Equilibria	11
4.1	Further Definitions	12
4.2	Discrete and Continuous Auctions: A Computational Equivalence	14
4.3	PPAD-completeness	17
5	Symmetric Equilibria in IID Auctions	18
6	Correlated Equilibria	25
7	Conclusion and Future Work	27
	Appendix	28
A	A Brief Discussion on Equilibrium Notions	28
B	Non-Existence of Approximate Pure Equilibria	28
C	NP-Hardness of Deciding Approximate Equilibria	32
D	PPAD-Hardness of Finding Mixed Equilibria	41
E	Existence of Symmetric Equilibria in IID Auctions	47

1 Introduction

The first-price auction is arguably the simplest and most intuitive auction format: there is one good for sale, which gets allocated to the highest bidder (breaking ties uniformly at random), for a payment equal to her submitted bid. In its primitive forms, this type of auction dates back to ancient times [Cas67], and became quite popular during the 17th century for the sale of paintings and estates. In recent times, the first-price auction has been central to the sale of online advertising space, known as the multi-billion-dollar industry of “ad impressions” (e.g., see [BDMMZ21])¹. Indeed, most digital marketplaces, known as *ad exchanges*—including Google’s Ad Manager—are currently using first-price auctions for the sale of their ad impressions [PST20; DRS21; CKSS22]. This mass migration to the first-price auction from other auction formats, such as the celebrated second-price auction [Vic61], was coined “the first-price movement” and was mainly attributed to the auction’s simplicity for the users, but also to increased revenue for the publishers. Indeed, in a survey of 100 publishers in November 2018, 78% reported an increase in revenue due to migrating to the first-price auction [Dig19].

Contrary to its famous second-price counterpart, the first-price auction provides incentives to the participants to strategically choose their bids. For example, consider a bidder that is willing to pay \$10 for the good. If the bidder believes that it is unlikely for any other bidder to be willing to bid an amount equal or close to \$10, she would *underbid*, aiming to win the good at a lower price. This induces a *strategic game*, in which the bidders choose their bids aiming to maximize their personal utilities. Inherent in this game is the notion of a *belief* of a bidder for the values (the aforementioned “willingness to buy”) of the other bidders, and hence, in game-theoretic terms, this is a *Bayesian game* of incomplete information. These games were formally introduced in the seminal work of Harsanyi [Har67] and capture situations in which the bidders are uncertain about the type of opponents they are facing in the game. The stable outcomes of these games, where no bidder wishes to unilaterally change her strategy, are known as *Bayes-Nash equilibria* (or simply, *equilibria*).

In the context of auctions, incomplete information is typically captured by probability distributions (also known as *priors*) over the possible values that the other bidders have for the item. Interestingly, Vickrey, in his seminal paper in 1961, defined and analysed a Bayesian game for the first-price auction, even before Harsanyi’s formal introduction of the concept. Following his work, a plethora of papers in the economics literature studied the equilibria of first-price auction; see, e.g., [Vic61; GLS67; MR85; MR00; MR03; Chw89; Plu92; Leb96; Leb99; Leb06; LP00; Ath01]. The focus of these works is primarily on showing the existence of equilibria for different variants of the auction, their (non)-uniqueness, and in relatively few cases how to “find” them, i.e., deriving closed-form solutions that characterize them.

In the computer science community, Filos-Ratsikas et al. [FGHLP23] recently put forward the study of the *computational complexity* of finding these equilibria; that is, devising polynomial-time algorithms to compute them, or proving computational hardness results for appropriate complexity classes. They studied a setting with (subjective) continuous priors and discrete bids, for which the existence of a (pure) equilibrium is guaranteed by the work of Athey [Ath01]. Among other results, they showed that computing an equilibrium of the auction is PPAD-complete in general, but polynomial-time solvable in special cases. The class PPAD was defined by Papadimitriou [Pap94] and is believed to contain problems that are computationally hard to solve, including finding a mixed Nash equilibrium in general games [DGP09; CDT09], among many others. Following up on [FGHLP23], Chen and Peng [CP23] considered the same setting with *common* priors and provided a PPAD-completeness result, but under a different tie-breaking rule, rather than uniform tie-breaking.

A common feature of the aforementioned works is that they assume the priors to be *continuous*, meaning that the values are drawn from continuous distributions over an infinite domain. While this assumption is quite common in the auction literature, considering auctions with *discrete* values has

¹In fact, the sale of ad impressions accounted for 58% of Google’s revenue in 2022, for an amount of \$162.45 billion; see [Sta24].

clear merits as well. For example, such priors are often formed based on past bidding behaviour, or can be interpreted as the prices of reselling the good²; in either case, they are inherently of a discrete nature [EMS09]. Additionally, in experiments designed to evaluate the incentive properties of the auction, the values are ultimately chosen from a discrete set, e.g., see [CST80; CRS82; CSW83; CSW88; Har89]. Furthermore, cryptographic implementations of the auction, which are based on *perturbations* of the values, are contingent on the discreteness in order to reason about equilibria in the perturbed space [MNT09]. Finally, the setting of discrete values and discrete bids makes the problem more amenable to computational approaches, as the representation of the inputs (the priors) and the outputs (the strategies) is straightforward. In contrast, the issue of representation in continuous value spaces is a rather intricate one, and requires additional monotonicity assumptions on the strategies; see the discussion in [FGHLP23] for more details. Other works that consider discrete priors in auctions include [CDW12; FHHK14; PP15; KBP18; ACLS22].

Motivated by the above, we consider the problem of computing equilibria in first-price auctions with discrete values and discrete bids. We phrase the following general question:

Open Question 1. What is the computational complexity of finding equilibria in first-price auctions with discrete values and discrete bids?

1.1 Our Results and Discussion

1.1.1 Pure equilibria

Contrary to the case of continuous values, it is known that in a first-price auction with discrete values, *pure* Bayes-Nash equilibria may not even exist [MR85; EMS09]. This motivates the computational question of *deciding* their existence. To this end, we provide our first main result:

Informal Theorem 1. *The problem of deciding the existence of a pure Bayes-Nash equilibrium, in a first-price auction with discrete subjective priors and discrete bids, is NP-complete.*

The term “subjective” in the statement of [Informal Theorem 1](#) refers to the fact that different bidders may have different beliefs about the values of other bidders. Subjective priors are also present in most of the results in [FGHLP23], as well as in the definition of general Bayesian games (e.g., see [Har67; Mye97; JR01]). Related notions of subjectivity have been considered rather widely in the literature, e.g., see [Hah73; FL86; BGM92; BG97; KL93; RW94; WP12; FW16]. An interesting special case of subjective priors is the well-known *independent private values (IPV)* model, where the prior distribution for every bidder i is the same from the perspective of any other bidder j (see, e.g., [MR85]).

ϵ -approximate Equilibria

In the context of game theory, it is often useful to talk about ϵ -approximate equilibria, i.e., sets of strategies where any unilateral deviation could improve a player’s utility by at most ϵ (here, in the additive sense). With regard to computational complexity in particular, this might be necessary, as exact equilibria in certain settings might involve irrational quantities in their description, making them a priori impossible to represent on a computer; see [FGHLP23] for a discussion related to the first-price auction, and also [Gol11] for a more general context. In our setting with discrete values and bids, this is not an issue: all pure equilibria are rational. Still, in light of the aforementioned non-existence and NP-hardness results, one might wonder: “Could ϵ -approximate equilibria always exist, and could we perhaps compute them in polynomial time?”. It turns out that the answer to both of these questions is “no”, at least when ϵ is a reasonably small constant. In fact, the formal version of [Informal Theorem 1](#)

²Experimental works in auctions, e.g. [CST80; CSW88], refer to the values explicitly as “resale values” when communicating them to the participants.

is an *inapproximability* result for ε -approximate equilibria, which we present after we establish a strengthening of the non-existence results of [MR85; EMS09] to approximate equilibria as well.

1.1.2 Mixed Equilibria

Motivated by the non-existence of pure equilibria for discrete values, we then turn our attention to *mixed* Bayes-Nash equilibria of the first-price auction, in which agents are allowed to randomize among different bids; note that previous works [FGHLP23; CP23] focus exclusively on the complexity of pure equilibria. The existence of such mixed equilibria is guaranteed, due to known existence results in general Bayesian games; see also our discussion in Section 6. This renders the (mixed) equilibrium-finding problem a *total* search problem; i.e., one that always has a solution, and thus a candidate for membership in the aforementioned class PPAD. Our second main result establishes that the problem is, in fact, PPAD-complete.

Informal Theorem 2. *The problem of computing a mixed Bayes-Nash equilibrium of a first-price auction with discrete subjective priors and discrete bids is PPAD-complete.*

We remark that, similar to Informal Theorem 1, the PPAD-hardness part of Informal Theorem 2 also applies even to ε -approximate equilibria, for a constant ε , whereas the PPAD-membership applies to approximate equilibria for the level of accuracy required by the representation (i.e., inversely-exponential in the input size).

One could conceivably exploit the nature of the aforementioned general existence proofs, and combine them with the general membership technique of Daskalakis, Fabrikant, and Papadimitriou [DFP06], to obtain the PPAD-membership part of our result in Informal Theorem 2. Here we take a different route, by establishing an interesting connection between the discrete and continuous auction settings. In particular, we prove a *computational equivalence* between ε -approximate *mixed* equilibria of the discrete first-price auction and ε' -approximate *pure* equilibria of the continuous variant. This equivalence could serve as a useful tool in the future: the gist of our framework is that one can effectively think of pure equilibria in the continuous setting and mixed equilibria in the discrete setting as two sides of the same coin, allowing for translating results between the two settings.

For example, in this paper we will use this equivalence to derive the aforementioned PPAD-membership as a corollary of the corresponding result for pure equilibria [FGHLP23], by transforming the discrete auction into a continuous one, and computing a pure equilibrium there. The other direction of the equivalence yields a PPAD-hardness result, which is however somewhat unsatisfactory, as it only applies to *monotone* equilibria. This is a by-product of the fact that, as mentioned earlier, pure equilibria of the continuous auction need to be monotone for reasons related to their representation [FGHLP23]. For that reason, we provide a stand-alone, stronger PPAD-hardness result for all equilibria (not necessarily monotone) via a direct reduction from the PPAD-complete problem PURE CIRCUIT of Deligkas et al. [DFHM22]. This allows also our PPAD-hardness to hold even for instances with very simple bidding spaces.

1.1.3 Positive Results

The NP- and PPAD-hardness results of the aforementioned theorems showcase the computational challenges in the quest for finding equilibria for general prior distributions. In quest for positive results, we first consider a natural restriction of these priors, to be independent and identically distributed (iid), and the question of computing *symmetric* mixed Bayes-Nash equilibria, i.e., equilibria in which all the bidders use the same bidding strategy. For this case we provide a PTAS (polynomial-time approximation scheme); that is, an algorithm that finds an ε -approximate equilibrium for any ε , in time polynomial in the description of the auction, but possibly exponential in $1/\varepsilon$. In particular, if we aim for

an ε -approximate equilibrium with ε constant, then this yields a polynomial-time algorithm. We state the corresponding theorem informally below:

Informal Theorem 3. *The problem of computing a symmetric ε -approximate mixed Bayes-Nash equilibrium of a first-price auction with iid priors admits a PTAS.*

Our proof is based on formulating the equilibrium computation problem as a system of polynomial inequalities, which can be solved within inversely-exponential precision using known results from the literature (e.g., see [GV88]). To ensure that the running time of our algorithm is polynomial, we need to make sure that the number of variables in the system is small; in our case, that would correspond to the value and bidding spaces being of fixed size. For the former, we exploit the nature of monotone strategies and devise a succinct representation which we refer to as *support-representation*. For the latter, we prove a “shrinkage lemma”, which enables us to work on a substantially smaller bidding space of size $O(1/\varepsilon)$, find an ε' -approximate equilibrium there, and translate it to an $(\varepsilon + \varepsilon')$ -equilibrium of the original auction.

Finally, we consider *correlated* equilibria, a more general equilibrium notion due to Aumann [Aum74]. In Bayesian games, there are several conceivable definitions of correlated equilibria; see, e.g., [Aum87; BM13; BM16; For93; For06; For23] and [Mye97, Sec. 6.3]. In this paper we adopt a standard notion, defined via the *type-agent representation*,³ a canonical normal-form game representation of Bayesian games that preserves the underlying fundamental equilibrium structure of the auction (e.g., see [Mye97, Sec. 2.8], [JR01, Sec. 7.2.3]), also adopted in [HST15; AB24]. Our contribution here, which is presented in Section 6, is showing that (exact) correlated equilibria of the first-price auction can be computed in polynomial time.

1.2 Further Related Work

As we mentioned in the introduction, Filos-Ratsikas et al. [FGHLP23] were the first to study the complexity of equilibrium computation in first-price auctions with continuous priors and discrete bids, providing a PPAD-completeness result for the case of subjective priors. In follow-up work, Chen and Peng [CP23] considered the IPV setting and proved a PPAD-completeness under an alternative, somewhat involved tie-breaking rule for the possible winners, rather than the standard uniform tie-breaking rule. Our shrinkage lemma is inspired by a conceptually similar idea in [CP23], but, contrary to their work, it refers to auctions with discrete values (rather than continuous), subjective priors (rather than IPV), and mixed equilibria (rather than pure equilibria).

The setting with discrete values (and discrete bids), which we study in the present paper, was first studied by [EMS09], who provided preliminary results limited mainly to 2 bidders with bivalued iid distributions. In a conceptually related paper, Wang, Shen, and Zuo [WSZ20] considered the computation of equilibria in first-price auctions with discrete values and continuous bids under a non-standard tie-breaking rule, which was used primarily by Maskin and Riley [MR00] as an intermediate step to prove results for uniform tie-breaking. Other works that study the complexity of equilibrium computation in Bayesian auctions, and Bayesian games in general, include [GGM07; CS08; PR08; CP14].

2 Preliminaries

In a (discrete, Bayesian) *first-price auction (DFPA)*, there is a set $N = \{1, 2, \dots, n\}$ of *bidders* and one item for sale. Each bidder i has a *value* $v_i \in V_i$ for the item and submits a *bid* $b_i \in B$. Sets V_1, V_2, \dots, V_n, B are *finite* subsets of $[0, 1]$ and are called the *value spaces* of the bidders and the *bidding space* of the auction, respectively.

³Some works refer to this representation as the *agent normal form representation*, e.g., see [HST15; AB24]. We use the term “type-agent representation” instead, following Myerson [Mye97, Sec. 2.8].

The item is allocated to the highest bidder, who has to submit a payment equal to her bid. In case of a tie for the highest bid, the winner is determined according to the *uniform tie-breaking* rule. That is, for a *bid profile* $\mathbf{b} = (b_1, \dots, b_n)$, the *ex-post utility* of bidder i with value v_i is defined as

$$\tilde{u}_i(\mathbf{b}; v_i) := \begin{cases} \frac{1}{|W(\mathbf{b})|}(v_i - b_i), & \text{if } i \in W(\mathbf{b}), \\ 0, & \text{otherwise,} \end{cases} \quad \text{where } W(\mathbf{b}) = \operatorname{argmax}_{j \in N} b_j \quad (1)$$

Bayesian priors. Each bidder $i \in N$ has a *subjective belief* for the values of each of the other bidders $j \in N \setminus \{i\}$, in the form of a *prior* distribution $F_{i,j}$ over V_j . This induces a product distribution $F_{-i} := \times_{j \neq i} F_{i,j}$ for the values $\mathbf{v}_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n) \in \times_{j \in N \setminus \{i\}} V_j =: V_{-i}$ of the other bidders. In other words, from the perspective of bidder i , the values v_j for $j \neq i$ are drawn *independently* from distributions $F_{i,j}$.

We will also be interested in special cases of these Bayesian priors, namely:

- *Independent Private Values (IPV)*, where $F_{i,j} = F_{i',j}$ for all $j \in N$ and $i, i' \in N \setminus \{j\}$. In this case, notation can be simplified by using F_j instead of $F_{i,j}$. This can be interpreted as the value profile $\mathbf{v} \in \times_{i \in N} V_i$ being drawn from the product distribution $F = \times_{j \in N} F_j$.
- *Identical Independent Values (iid)*, which is a special case of IPV above where $V_i = V_{i'}$ and $F_i = F_{i'}$ for all $i, i' \in N$. In other words, bidder values are iid according to a common distribution F .

2.1 The Auction Game

The DFPA described above gives rise to a Bayesian game of incomplete information, where each bidder i chooses her bid based on her own (true) value v_i and her beliefs F_{-i} about the other bidders. A (mixed) *strategy* of bidder i is a function $\beta_i : V_i \rightarrow \Delta(B)$ mapping values to distributions over bids.^{4,5} *Pure* strategies correspond to the special case where a mixed strategy β_i always assigns full mass on single bids; that is, for all $v_i \in V_i$ there exists a $b_i \in B$ such that $\beta_i(v_i)(b_i) = 1$ and $\beta_i(v_i)(b) = 0$ for all $b \neq b_i$. Therefore, for simplicity, we will sometimes represent pure strategies directly as functions $\hat{\beta}_i : V_i \rightarrow B$ from values to bids.

Given a strategy profile $\beta_{-i} \in \times_{j \in N \setminus \{i\}} \Delta(B)^{V_j}$ of the other bidders, the (interim) *utility* of a bidder i with value v_i , when bidding $b \in B$, is given by

$$\begin{aligned} u_i(b, \beta_{-i}; v_i) &:= \mathbb{E}_{\mathbf{v}_{-i} \sim F_{-i}} \left[\mathbb{E}_{\mathbf{b}_{-i} \sim \beta_{-i}(\mathbf{v}_{-i})} [\tilde{u}_i(b, \mathbf{b}_{-i}; v_i)] \right] \\ &= \sum_{\mathbf{v}_{-i} \in V_{-i}} \left(\prod_{j \in N \setminus \{i\}} f_{i,j}(v_j) \right) \sum_{\mathbf{b}_{-i} \in B^{N \setminus \{i\}}} \left(\prod_{j \in N \setminus \{i\}} \beta_j(v_j)(b_j) \right) \tilde{u}_i(b, \mathbf{b}_{-i}; v_i). \end{aligned} \quad (2)$$

where $\beta_{-i}(\mathbf{v}_{-i})$ is a shorthand for the product distribution $\times_{j \in N \setminus \{i\}} \beta_j(v_j)$, and $f_{i,j}$ denotes the probability mass function (pmf) of $F_{i,j}$. Recall that $\beta_j(v_j)(b_j)$ denotes the probability that bidder $j \neq i$ submits bid b_j when having value v_j . This can be viewed as the bidder computing her utility as follows: (a) she draws a value v_j for each bidder $j \neq i$ from her corresponding subjective prior $F_{i,j}$; (b) she uses the strategy “rules” β_{-i} of the others to map the values obtained in (a) to actual bid distributions over B .

⁴See, e.g., [Mye97, Sec. 3.9], [JR01, Sec. 7.2.3], [MW85; LB17].

⁵Formally, for any finite set X we define $\Delta(X)$ as the simplex

$$\Delta(X) := \left\{ p \in [0, 1]^X \mid \sum_{x \in X} p(x) = 1 \right\},$$

where $p(x)$ can be interpreted as the mass that the distribution assigns to element $x \in X$. We will denote the support of a distribution $p \in \Delta(X)$ with $\operatorname{supp}(p) := \{x \in X \mid p(x) > 0\}$.

For convenience, we also extend, in the natural way, the definition above to handle the case where bidder i randomizes over her bids; that is, for $\boldsymbol{\gamma} \in \Delta(B)$ we define:

$$u_i(\boldsymbol{\gamma}, \boldsymbol{\beta}_{-i}; v_i) := \mathbb{E}_{b \sim \boldsymbol{\gamma}} [u_i(b, \boldsymbol{\beta}_{-i}; v_i)] = \sum_{b \in B} \boldsymbol{\gamma}(b) u_i(b, \boldsymbol{\beta}_{-i}; v_i). \quad (3)$$

Equilibria. We proceed to define the two main (exact and approximate) equilibrium notions that we will study in this paper.

Definition 1 (ε -approximate mixed Bayes-Nash equilibrium of the DFPA). Let $\varepsilon \geq 0$. A (mixed) strategy profile $\boldsymbol{\beta} = (\beta_1, \dots, \beta_n)$ is an (interim) ε -approximate mixed Bayes-Nash equilibrium (MBNE) of the DFPA if for any bidder $i \in N$ and any value $v_i \in V_i$,

$$u_i(\beta_i(v_i), \boldsymbol{\beta}_{-i}; v_i) \geq u_i(\boldsymbol{\gamma}, \boldsymbol{\beta}_{-i}; v_i) - \varepsilon \quad \text{for all } \boldsymbol{\gamma} \in \Delta(B). \quad (4)$$

We will refer to a 0-approximate MBNE as an *exact* MBNE.

In the special case where all bidders choose the same strategies, i.e. $V_i = V_{i'}$ and $\beta_i = \beta_{i'}$ for all $i, i' \in N$, the (approximate) equilibrium will be called *symmetric*.

Remark 1. It is straightforward to check that condition (4) can be equivalently stated to just range over all *pure* deviations $\boldsymbol{\gamma} \in B$ (instead of all *mixed* ones $\boldsymbol{\gamma} \in \Delta(B)$), without affecting [Definition 1](#).

It is known (see, e.g., [JR01, Theorem 7.3]) that finite Bayesian games (and thus, DFPA as well) always have at least one exact MBNE. Therefore, existence of ε -approximate MBNE is also guaranteed for any $\varepsilon > 0$.

Similarly, one can define the notion of pure equilibria:

Definition 2 (ε -approximate pure Bayes-Nash equilibrium of the DFPA). Let $\varepsilon \geq 0$. A pure strategy profile $\hat{\boldsymbol{\beta}} = (\hat{\beta}_1, \dots, \hat{\beta}_n)$ is an (interim) ε -approximate pure Bayes-Nash equilibrium (PBNE) of the DFPA if for any bidder $i \in N$ and any value $v_i \in V_i$,

$$u_i(\hat{\beta}_i(v_i), \hat{\boldsymbol{\beta}}_{-i}; v_i) \geq u_i(b, \hat{\boldsymbol{\beta}}_{-i}; v_i) - \varepsilon \quad \text{for all } b \in B.$$

In contrast to mixed equilibria, discrete first-price auctions do not, in general, have exact pure equilibria [EMS09]. In [Theorem 3.1](#), we extend this nonexistence result to *approximate* (pure) equilibria.

No overbidding. Following the literature of first-price auctions in economics [MR00; MR03; Leb06], as well as in computer science [FGHLP23; WSZ20; EMS09; CP23], we will also make the standard no overbidding assumption. In our context, this translates into strategies never assigning positive probability to some bid b_i which is larger than the bidder's value v_i . Note that, if we allow the bidders to abstain from the auction, by adopting what is known as the *null bid* assumption in the literature (see [MR85; Ath01]), i.e. that $0 \in B$, then any overbidding strategy is weakly dominated by a strategy that transfers any bidding probability mass from a $b_i > v_i$ to the zero bid.

Representation. For our model to be fit for computational purposes, we have to determine how exactly the inputs and the outputs of the associated computational problems will be represented. Similarly to [FGHLP23; CP23], we will assume that the bidding space B is given explicitly as part of the input in the form of rational numbers. The value spaces V_i will also be given explicitly, in the same way. The value distributions $F_{i,j}$ will be described explicitly via their probability mass functions. That is, we are given a list of rationals $\{f_{i,j}(v)\}_{v \in V_j}$ representing the probability that bidder j has true value v , from the perspective of bidder i . The mixed strategies in the output are described explicitly by (rational) numbers $\{p_i(v, b)\} \in [0, 1]$ representing the probability that bidder i submits bid $b \in B$ when her true value is $v \in V_i$. Note that, for the special case of pure strategies it is $\{p_i(v, b)\} \in \{0, 1\}$.

2.2 Efficient Computation of Bidder Utilities

We conclude our preliminaries with the following lemma, which will be useful for several of our results throughout the paper. The lemma establishes that the expected utilities, given mixed bidding strategies of the bidders, can be computed in polynomial time. We remark that Filos-Ratsikas et al. [FGHLP23, Lemma 3.2] proved a similar lemma for the case of pure strategies (and continuous value spaces). Our proof is similar, but extra care is needed in order to handle the case of mixed strategies; we use similar notation to [FGHLP23], for consistency.

Lemma 2.1. *Fix a DFPA. For any bidder $i \in N$ and any true value $v_i \in V_i$, the utility⁶ $u_i(\boldsymbol{\gamma}, \boldsymbol{\beta}_{-i}; v_i)$ of i , given as input any distribution $\boldsymbol{\gamma}$ over her bids and any mixed bidding strategy profile $\boldsymbol{\beta}_{-i}$ of the other bidders, is computable in polynomial time.*

Proof. Without loss of generality, we can reorder the bidders such that the one whose utility we are calculating is the last, i.e., $n = |N|$. Let $H_n(b, \boldsymbol{\beta}_{-n})$ denote the probability of bidder n winning when bidding $b \in B$ against the mixed strategies $\boldsymbol{\beta}_{-n}$ of the remaining bidders (based on bidder n 's subjective beliefs). We can write bidder n 's utility when bidding b as $u_n(b, \boldsymbol{\beta}_{-n}; v) = (v - b)H_n(b, \boldsymbol{\beta}_{-n})$, so it suffices to show how to efficiently compute $H_n(b, \boldsymbol{\beta}_{-n})$. We can express this as:

$$H_n(b, \boldsymbol{\beta}_{-n}) = \sum_{r=0}^{n-1} \frac{1}{r+1} T_n(b, n-1, r) \quad (5)$$

where, for $0 \leq r \leq \ell \leq n-1$, we use $T_n(b, \ell, r)$ to denote the probability that *exactly* r out of the first ℓ bidders bid b , and the remaining $\ell - r$ bidders all bid below b . Once again, this probability is based on bidder n 's subjective priors. Next, for a given bidder j , let:

$$g_{j,b} := \mathbb{E}_{v_j \sim F_{n,j}} \left[\Pr_{\xi \sim \beta_j(v_j)} [\xi = b] \right] = \sum_{v_j \in V_j} \beta_j(v_j)(b) f_{n,j}(v_j) \quad (6)$$

$$G_{j,b} := \mathbb{E}_{v_j \sim F_{n,j}} \left[\Pr_{\xi \sim \beta_j(v_j)} [\xi < b] \right] = \sum_{\substack{b' \in B, \\ b' < b}} g_{j,b'} \quad (7)$$

where $f_{n,j}$ is the probability mass function corresponding to bidder n 's beliefs about bidder j 's value. What we have defined above is the probability (perceived from the perspective of bidder n) that bidder j bids *exactly* b (denoted as $g_{j,b}$) and below b (denoted as $G_{j,b}$). The computation of these takes into account two sources of randomness; first, we draw a value v_j for bidder j from the prior distribution $F_{n,j}$, and then we pick a bid ξ , given the distribution over bids that j 's strategy represents at the drawn value $\beta_j(v_j)$. It is easy to see that these quantities can be computed in time $O(|V||B|)$ with access to the pmf of the prior distributions as well as the distributions over bids corresponding to the mixed strategies. We can now express the term $T_n(b, n-1, r)$ from (5) using the newly defined quantities of Equations (6) and (7):

$$T_n(b, n-1, r) = \sum_{\substack{S \subseteq [n-1] \\ |S|=r}} \prod_{j \in S} g_{j,b} \cdot \prod_{j \notin S} G_{j,b} \quad (8)$$

where we have used the notation $[n-1]$ to indicate the set of the integers from 1 to $n-1$ inclusive.

Note that (8) cannot be computed efficiently in any obvious way, since the number of summands can be exponential in n . To overcome this issue, we proceed with a dynamic programming algorithm for computing $T_n(b, n-1, r)$, in a similar manner to [FGHLP23, Lemma 3.2]. The probabilities $T_n(b, \ell, k)$

⁶See (3) and (2) for the utility definition.

can be computed from $g_{j,b}$ and $G_{j,b}$ (which were defined in [Equations \(6\) and \(7\)](#) respectively), via dynamic programming, conditioning on bidder ℓ 's bid in the following way:

$$\begin{aligned} T_n(b, 0, 0) &= 1; \\ T_n(b, \ell, k) &= 0, && \text{for } k > \ell; \\ T_n(b, \ell + 1, 0) &= T_n(b, \ell, 0)G_{\ell+1,b}; \\ T_n(b, \ell + 1, k + 1) &= T_n(b, \ell, k)g_{\ell+1,b} + T_n(b, \ell, k + 1)G_{\ell+1,b}; && \text{for } k \leq \ell. \end{aligned}$$

Thus, all values of $T_n(b, n - 1, k)$, for $k = 0, \dots, n - 1$, can be computed with a total number of $O(n^2)$ recursive calls, so that $H_n(b, \boldsymbol{\beta}_{-n})$ can be computed in polynomial time. This directly implies that the expected utility $u_n(\boldsymbol{\beta}; v)$ of bidder n , given as input a mixed strategy profile $\boldsymbol{\beta}$, is efficiently computable, by virtue of it being a convex combination of at most $|B|$ pure strategy utilities. \square

3 Pure Equilibria

In this section, we present our results for the pure Bayes-Nash equilibria (PBNE) of first-price auctions. First, it is already known from [\[EMS09\]](#) that in the case of discrete values, PBNE may not exist, even for very simple instances of two bidders and iid priors with support size 2. In the following theorem, we strengthen this impossibility result to ε -approximate PBNE, for a reasonably large constant value of ε . We delegate the proof of the theorem to [Appendix B](#).

Theorem 3.1. *For any $\varepsilon < \frac{1}{72}$, there exist discrete first-price auctions, even with only two bidders and iid priors with support size 2, that do not have ε -approximate pure equilibria.*

Driven by the non-existence result of [Theorem 3.1](#), we consider the computational complexity of deciding the existence of PBNE in a DFPA. Our main result of this section is that, when the prior distributions are subjective, the problem is NP-complete, even for approximate equilibria. The following statement is the formal version of [Informal Theorem 1](#). In the interest of space and the flow of our presentation, we provide only a high-level sketch here. The full proof of our theorem is delegated to [Appendix C](#).

Theorem 3.2. *Consider any $\varepsilon < \frac{1}{180}$. The problem of deciding the existence of an ε -PBNE of a DFPA is NP-complete.*

Proof sketch. We begin by showing membership in NP, demonstrating that we can verify any positive certificate in polynomial time. Given a pure strategy profile $\hat{\boldsymbol{\beta}}$, it suffices to efficiently verify that strategy $\hat{\beta}_i$ is an ε -best-response for each bidder $i \in N$ and value $v_i \in V_i$. To this end, we can quantify over all, polynomially many in the size of the input, possible values of $i \in N, v_i \in V_i, b \in B$, and compute the corresponding utilities. By [Lemma 2.1](#), this can be done in polynomial time.

We now turn to the NP-hardness of the problem. We reduce from the well-known NP-complete Circuit Satisfiability problem (CIRCUIT-SAT) [\[GJ79\]](#). For ease of notation, our reduction will use the bidding space $B = \{0, 1, 2, 3\}$ and the value space (common for all bidders) $V = \{0, \frac{9}{4}, 10\}$. With this setup, we will prove the result for $\varepsilon < 1/18$. Dividing everything by the largest entry (in this case 10), we bring the problem back to the desired domain $[0, 1]$, while keeping all comparisons between utilities in the proof intact. Hence, we obtain the result for our version of the problem for $\varepsilon < 1/180$.

The proof can be outlined in the following steps:

- We define the set of valid strategies (mapping from values to bids) $\{s_0, s_1\}$ that correspond to the boolean values *false* and *true*, respectively.

- We introduce 3 gadgets (*projection*, *OR*, and *NOT*), which are used to simulate the *OR* and *NOT* gates of the circuit, as well as a *split node*, which is used to copy the value of its input to two outputs. Note that any boolean circuit can be converted in polynomial time to another one that only uses these gates. The *projection* gadget is used in order to correctly map any bidder that plays a strategy different than s_0 and s_1 back to the set of valid strategies. By introducing bidders with appropriate subjective prior distributions, we establish that all the induced bidders can play ε -best-responses simultaneously if and only if the boolean values specified by their strategies satisfy the definitions of the corresponding gates of the circuit.
- For the simulation step, we assumed that all input bidders played a valid strategy. To ensure this, we introduce the *input gadget*. This can be seen as introducing an extra bidder i' for each bidder i corresponding to an input of the circuit, with subjective priors resembling a *projection* gate from i to i' and vice versa. We show in our analysis that this guarantees that i and i' can simultaneously play ε -best-responses if and only if they both play the same, valid, strategy.
- To finish off the construction, it remains to encompass the notion of the satisfiability of the circuit, meaning distinguishing the cases where the output can get the value *true*. To achieve this, we introduce the final gadget of our reduction, the *output gadget*. This introduces two new bidders to the DFPA, k and ℓ , which initially only have subjective beliefs that put positive probability mass to some value greater than 0 for each other. We show that they cannot simultaneously play ε -best-responses. Then, we describe k 's prior distribution for the bidder i corresponding to the output of the circuit, such that in the resulting DFPA it is possible for all 3 bidders (i, k, ℓ) to simultaneously play ε -best-responses if and only if i plays strategy s_1 . Therefore, combining this with the previous steps, the resulting DFPA has an ε -PBNE for $\varepsilon < \frac{1}{18}$ if and only if the circuit is satisfiable. \square

4 Mixed Equilibria

In this section, we switch our focus to the more general class of mixed equilibria. At first glance, one might be inclined to believe that since both the value space and the bidding space are discrete, the existence of MBNE follows immediately by Nash's Theorem [Nas51]. However, this is not the case, since the DFPA game is a Bayesian game of incomplete information. MBNE in (general) Bayesian games have been extensively discussed in the literature (see, e.g., [MW85; Aum61], [JR01, Chapter 7.2.3], and [Mye97, Chapter 2.8]). A common way to prove their existence is via transforming the Bayesian game into an appropriately constructed normal form game, and invoking Nash's Theorem there; see, e.g., [JR01, Theorem 7.3]. This normal form game formulation is the *type-agent representation* that we mentioned in our introduction, and will be critical for our discussion on correlated equilibria later, in Section 6. However, it is not needed for any of our results of the present section. Note that these general existence results for mixed equilibria in Bayesian games directly imply the existence of a MBNE in DFPA as well, rendering the problem of finding a MBNE, a total search problem.

We will also provide the proof of Informal Theorem 2, namely the PPAD-completeness of computing MBNE of the DFPA. En route to this result, we will present an interesting *computational equivalence theorem* between approximate mixed Bayes-Nash equilibria of the DFPA with discrete values and approximate pure Bayes-Nash equilibria of the DFPA with continuous values, the setting studied in [FGHLP23; CP23]. This equivalence allows us to translate positive and negative results in either direction; importantly, for the purpose of this section, it will help us establish the PPAD-membership of computing a MBNE of the DFPA. In principle, the equivalence theorem could also be used to get a PPAD-hardness result for the same problem, albeit somewhat weak in nature. Instead, we will provide a stronger, stand-alone PPAD-hardness result for computing a MBNE, which does not rely on the equivalence theorem.

4.1 Further Definitions

Before we present our equivalence theorem, we need to introduce some necessary terminology for the setting of first-price auctions with continuous priors (and discrete bids), as well as some further definitions that will be useful for the subsequent results.

Continuous first-price auctions. A continuous first-price auction (CFPA) is defined very similarly to DFPA, the key difference being that the value spaces V_i are continuous intervals of real numbers and the prior distributions $F_{i,j}$ are continuous. The bidding space $B = \{b_1, \dots, b_m\}$ is still discrete. For the purpose of this paper, it is enough to consider distributions with piecewise constant densities, represented explicitly by the endpoints and the height of the intervals. As we mentioned in the introduction, the representation of strategies in a CFPA is not as straightforward as in the case of a DFPA. For this reason, one needs to focus on *monotone strategies*, meaning that bidders submit higher bids on higher values. Using this, a strategy can be represented as a sequence of *jump points* $\alpha_1, \dots, \alpha_{m-1}$, where $\alpha_i \in V_i$ is the largest value for which bidder i bids b_i or lower. The notion of a pure Bayes-Nash equilibrium (PBNE) of a CFPA is defined similarly to a DFPA. We refer the reader to [FGHLP23, Sec. 2] for more details.

It will be useful to define a notion of monotonicity, appropriate for mixed strategies:

Definition 3 (Monotone mixed strategies in a DFPA). Consider a DFPA with value space V_i for each bidder i . A mixed strategy $\beta_i \in \Delta(B)^{V_i}$ (of bidder i) will be called *monotone* if

$$\max \text{supp}(\beta_i(v)) \leq \min \text{supp}(\beta_i(v')) \quad \text{for all } v, v' \in V_i \text{ with } v < v'.$$

A strategy profile will be called monotone, if the strategies of all bidders in it are monotone.

Next, we present the notion of ε -well-supported MBNE of a DFPA, a more refined notion of equilibrium approximation which we will use in the results of this section.

Definition 4 (ε -well-supported mixed Bayes-Nash equilibrium of the DFPA). Let $\varepsilon \geq 0$. A (mixed) strategy profile $\beta = (\beta_1, \dots, \beta_n)$ is an (interim) ε -well-supported mixed Bayes-Nash equilibrium (MBNE) of the DFPA if for any $i \in N$, any $v_i \in V_i$, and any $b \in \text{supp}(\beta_i)$:

$$u_i(b, \beta_{-i}; v_i) \geq u_i(\gamma, \beta_{-i}; v_i) - \varepsilon \quad \text{for all } \gamma \in \Delta(B). \quad (9)$$

Informally, in an ε -well-supported MBNE, the support of a strategy only contains bids that approximately (within ε) maximize the bidder's utility.⁷

Finally, we define a technical notion, which will be useful for us in the following:

Definition 5 (Interaction degree of a DFPA). A DFPA has *interaction degree* bounded by $d \in \{1, \dots, n-1\}$, if for any bidder $i \in N$, there exists a set $J \subseteq N \setminus \{i\}$ of bidders with $|J| \leq d$ such that for any bidder $j \in N \setminus (J \cup \{i\})$ we have $0 \in V_j$ and $f_{i,j}(0) = 1$.

Note that any DFPA has interaction degree bounded by $n - 1$. Intuitively, the interaction degree provides a bound on how many bidders are perceived as always having zero value, from the perspective of some other bidder. The following lemma shows that, as long as the interaction degree is constant, ε -well-supported MBNE can be obtained in polynomial time from ε' -approximate MBNE of the DFPA, when both ε and ε' are constant.

Lemma 4.1. *Consider any DFPA with interaction degree bounded by $d \in \{1, \dots, n-1\}$. Then, for any $\varepsilon \in [0, 1]$, a (monotone) ε -well-supported MBNE can be obtained in polynomial time given any (monotone) $\varepsilon^2/(8d)$ -approximate MBNE. Furthermore, if the auction is iid, then this also holds if we restrict our attention to symmetric equilibria.*

⁷Chen and Peng [CP23] use the term “ ε -well-supported” to refer to *interim* equilibria, see our Definition 1. To avoid any misconceptions, we provide a short discussion in Appendix A.

Proof. Let β be a δ -approximate MBNE of the auction, for some $\delta > 0$. Consider any $\gamma \geq \delta$, to be fixed later. For any $i \in N$ and any $v_i \in V_i$, let

$$G := \{b \in B : u_i(b, \beta_{-i}; v_i) \geq \max_{b' \in B} u_i(b', \beta_{-i}; v_i) - \gamma\}$$

be the set of bids that yield optimal utility up to γ . Then, we must have

$$\beta_i(v_i)(B \setminus G) := \sum_{b \in B \setminus G} \beta_i(v_i)(b) < \delta/\gamma \quad (10)$$

Indeed, if we had $\beta_i(v_i)(B \setminus G) \geq \delta/\gamma$, then

$$\begin{aligned} u_i(\beta_i(v_i), \beta_{-i}; v_i) &= \sum_{b \in B \setminus G} \beta_i(v_i)(b) \cdot u_i(b, \beta_{-i}; v_i) + \sum_{b \in G} \beta_i(v_i)(b) \cdot u_i(b, \beta_{-i}; v_i) \\ &< \beta_i(v_i)(B \setminus G) \cdot \left(\max_{b' \in B} u_i(b', \beta_{-i}; v_i) - \gamma \right) + \beta_i(v_i)(G) \cdot \max_{b' \in B} u_i(b', \beta_{-i}; v_i) \\ &= \max_{b' \in B} u_i(b', \beta_{-i}; v_i) - \gamma \cdot \beta_i(v_i)(B \setminus G) \\ &\leq \max_{b' \in B} u_i(b', \beta_{-i}; v_i) - \delta \end{aligned}$$

contradicting the fact that β is a δ -approximate MBNE.

Next, we construct a strategy profile β' from β . For every $i \in N$ and every $v_i \in V_i$ we proceed as follows. For any $b \in B \setminus G$ we set

$$\beta'_i(v_i)(b) := 0$$

and for any $b \in G$ we set

$$\beta'_i(v_i)(b) := \frac{\beta_i(v_i)(b)}{\sum_{b' \in G} \beta_i(v_i)(b')}$$

where the denominator is strictly positive by (10). Thus, $\beta'_i(v_i)$ is a well-defined distribution over B that only puts positive mass on bids in G . Since $\beta_i(v_i)$ is non-overbidding, so is $\beta'_i(v_i)$. Furthermore, we have $\text{supp}(\beta'_i(v_i)) \subseteq \text{supp}(\beta_i(v_i))$, and thus β' is monotone, if β is monotone. If the auction is iid and β is symmetric, then β' will also be symmetric.

Now, we have that for all $i \in N$ and $v_i \in V_i$, the total variation distance between $\beta_i(v_i)$ and $\beta'_i(v_i)$ satisfies

$$\text{TV}(\beta_i(v_i), \beta'_i(v_i)) := \frac{1}{2} \|\beta_i(v_i) - \beta'_i(v_i)\|_1 \leq \frac{1}{2} \cdot 2 \cdot \beta_i(v_i)(B \setminus G) < \delta/\gamma$$

by (10). Fix some $i \in N$. Let D_j and D'_j denote the distributions over B obtained by drawing $v_j \in V_j$ according to $F_{i,j}$, and then b according to $\beta_j(v_j)$, respectively $\beta'_j(v_j)$. We have

$$\begin{aligned} \text{TV}(D_j, D'_j) &= \frac{1}{2} \sum_{b \in B} \left| \sum_{v_j \in V_j} f_{i,j}(v_j) \left(\beta_j(v_j)(b) - \beta'_j(v_j)(b) \right) \right| \\ &\leq \sum_{v_j \in V_j} f_{i,j}(v_j) \frac{1}{2} \sum_{b \in B} |\beta_j(v_j)(b) - \beta'_j(v_j)(b)| \\ &\leq \sum_{v_j \in V_j} f_{i,j}(v_j) \delta/\gamma = \delta/\gamma. \end{aligned}$$

Next, define the product distributions $D := \times_{j \in N \setminus \{i\}} D_j$ and $D' := \times_{j \in N \setminus \{i\}} D'_j$ over B^{n-1} . Since the auction has interaction degree bounded by d , there exists a set $J \subseteq N \setminus \{i\}$ with $|J| \leq d$ such that for all

$j \in N \setminus (J \cup \{i\})$ we have $0 \in V_j$ and $f_{i,j}(0) = 1$. By the no-overbidding assumption, it must be that for all $j \in N \setminus (J \cup \{i\})$, $\beta_j(0)(0) = 1 = \beta'_j(0)(0)$, and thus $D_j = D'_j$. As a result

$$\text{TV}(D, D') \leq \sum_{j \in N \setminus \{i\}} \text{TV}(D_j, D'_j) = \sum_{j \in J} \text{TV}(D_j, D'_j) \leq d \cdot \delta / \gamma.$$

Now, fix some bid $b \in B$. For any $k \in \{0, 1, \dots, n-1\}$, let $E_k \subseteq B^{n-1}$ denote the event that exactly k bids are equal to b and the rest are strictly smaller. Recall the definition of the H functions in [Lemma 2.1](#). Then, we can write

$$\begin{aligned} H_i(b, \beta_{-i}) &= \sum_{k=0}^{n-1} \frac{1}{k+1} \Pr_D[E_k] \\ &= \frac{1}{n} \sum_{k=0}^{n-1} \Pr_D[E_k] + \sum_{\ell=1}^{n-1} \left(\frac{1}{\ell} - \frac{1}{\ell+1} \right) \sum_{k=0}^{\ell-1} \Pr_D[E_k] \\ &= \frac{1}{n} \Pr_D \left[\bigcup_{k=0}^{n-1} E_k \right] + \sum_{\ell=1}^{n-1} \left(\frac{1}{\ell} - \frac{1}{\ell+1} \right) \Pr_D \left[\bigcup_{k=0}^{\ell-1} E_k \right] \end{aligned}$$

where we used the fact that the events E_k are all disjoint. We also obtain the analogous expression for $H_i(b, \beta'_{-i})$ and D' . It follows that

$$\begin{aligned} &|H_i(b, \beta_{-i}) - H_i(b, \beta'_{-i})| \\ &\leq \frac{1}{n} \left| \Pr_D \left[\bigcup_{k=0}^{n-1} E_k \right] - \Pr_{D'} \left[\bigcup_{k=0}^{n-1} E_k \right] \right| + \sum_{\ell=1}^{n-1} \left(\frac{1}{\ell} - \frac{1}{\ell+1} \right) \left| \Pr_D \left[\bigcup_{k=0}^{\ell-1} E_k \right] - \Pr_{D'} \left[\bigcup_{k=0}^{\ell-1} E_k \right] \right| \\ &\leq \frac{1}{n} d\delta / \gamma + \sum_{\ell=1}^{n-1} \left(\frac{1}{\ell} - \frac{1}{\ell+1} \right) d\delta / \gamma = d\delta / \gamma \end{aligned}$$

where we used $\text{TV}(D, D') \leq d\delta / \gamma$.

As a result, for any $i \in N$, $v_i \in V_i$, and $b \in B$ we have

$$|u_i(b, \beta_{-i}; v_i) - u_i(b, \beta'_{-i}; v_i)| = |H_i(b, \beta_{-i}) \cdot (v_i - b) - H_i(b, \beta'_{-i}) \cdot (v_i - b)| \leq d\delta / \gamma.$$

Finally, for any $i \in N$, $v_i \in V_i$, and $b \in \text{supp}(\beta'_i(v_i)) \subseteq G$ we have

$$u_i(b, \beta_{-i}; v_i) \geq \max_{b' \in B} u_i(b', \beta_{-i}; v_i) - \gamma$$

by construction of G , and as a result

$$u_i(b, \beta'_{-i}; v_i) \geq \max_{b' \in B} u_i(b', \beta'_{-i}; v_i) - \gamma - 2d\delta / \gamma.$$

Thus, β' is a $(\gamma + 2d\delta / \gamma)$ -well supported MBNE. Setting $\gamma := \sqrt{2d\delta}$ and $\delta := \varepsilon^2 / (8d)$ yields the result. Note that for $\varepsilon = 0$ the statement trivially holds. \square

4.2 Discrete and Continuous Auctions: A Computational Equivalence

We are now ready to prove our computational equivalence result. This is composed of [Lemmas 4.2](#) and [4.3](#) below.

Lemma 4.2. *Given $\delta \in (0, 1)$ and a DFPA, we can construct in polynomial time a CFPA (with the same bidding space) such that for any $\varepsilon \geq 0$, we can transform in polynomial time any ε -PBNE of the CFPA into an $(\varepsilon + \delta)$ -well-supported monotone MBNE of the DFPA. Furthermore, this reduction maps IPV (resp. iid) auctions to IPV (resp. iid) auctions, and symmetric equilibria to symmetric equilibria.*

Proof. Let $\delta \in (0, 1)$ be given. Consider a DFPA given by the value spaces V_1, \dots, V_n and the distributions $F_{i,j}$. Without loss of generality, we can assume that δ satisfies

$$2\delta \leq \min_{p,q \in V_1 \cup \dots \cup V_n \cup B: p \neq q} |p - q|. \quad (11)$$

Indeed, if this is not the case, we can replace δ by a smaller value in $(0, 1)$ that satisfies this inequality.

We construct a CFPA with the same bidding space B and with distributions $F'_{i,j}$. For every $i \neq j$, the piecewise constant density function of the distribution $F'_{i,j}$ is constructed from $F_{i,j}$ as follows.

- For every $v_j \in V_j$ with $v_j \neq 0$, we let the pdf have value $f_{i,j}(v_j)/\delta$ in interval $[v_j - \delta, v_j]$.
- If $0 \in V_j$, we let the pdf have value $f_{i,j}(0)/\delta$ in interval $[0, \delta]$.
- The pdf has value 0 everywhere else in $[0, 1]$.

For an illustration, see [Figure 1](#). This pdf is well-defined by (11) (also using the fact that $0 \in B$), and in particular the blocks of value do not overlap. Furthermore, any additional structure of the auction (namely IPV or iid) is kept intact.

Now consider any ε -PBNE β' of the CFPA. We construct a corresponding strategy profile β in the original DFPA as follows. For any $j \in N$, $\beta_j : V_j \rightarrow \Delta(B)$ is given by:

- For every $v_j \in V_j$ with $v_j \neq 0$, let $\beta_j(v_j)$ be the distribution of $\beta'_j(u)$ where u is drawn uniformly at random in $[v_j - \delta, v_j]$. In other words, for all $b \in B$

$$\beta_j(v_j)(b) = \lambda(\{u \in [v_j - \delta, v_j] : \beta'_j(u) = b\})/\delta$$

where λ denotes the Lebesgue measure (in this case just the length of the interval).

- If $0 \in V_j$, let $\beta_j(0)$ be the distribution of $\beta'_j(u)$ where u is drawn uniformly at random in $[0, \delta]$.

Recall that $\beta'_j : [0, 1] \rightarrow B$ is a non-overbidding non-decreasing step-function strategy. As a result, and also by (11), the mixed strategy β_j will also be non-overbidding. Furthermore, the fact that β'_j is non-decreasing immediately yields that β_j is monotone. Finally, if $V_i = V_{i'}$ for all i, i' and β' is symmetric, then so is β .

Next, we note that the construction of $F'_{i,j}$ from $F_{i,j}$ and the construction of β from β' have been carefully devised to ensure that for every $i \neq j$ the two following distributions over B coincide

- Pick $v_j \in V_j$ according to $F_{i,j}$, and then output $b \in B$ according to $\beta_j(v_j)$.
- Pick $v \in [0, 1]$ according to $F'_{i,j}$ and output $b = \beta'_j(v)$.

As a result, we have that

$$u_i(b, \beta_{-i}; v_i) = u'_i(b, \beta'_{-i}; v_i) \quad (12)$$

for all $i \in N$, $b \in B$, and $v_i \in V_i$, where u_i, u'_i denote the utilities in the DFPA and CFPA respectively.

It remains to prove that β is an $(\varepsilon + \delta)$ -well-supported monotone MBNE of the DFPA. For this, it suffices to show that for any $i \in N$ and $v_i \in V_i$, we have that all $b \in B$ with $\beta_i(v_i)(b) > 0$ satisfy $u_i(b, \beta_{-i}; v_i) \geq u_i(b', \beta_{-i}; v_i) - \varepsilon - \delta$ for all $b' \in B$. Consider any such i, v_i, b and note that by construction of β_i from β'_i there exists $v \in [0, 1]$ with $|v - v_i| \leq \delta$ such that $\beta'_i(v) = b$. Since β' is an ε -PBNE β' of the CFPA, it follows that for all $b' \in B$

$$\begin{aligned} & u'_i(b, \beta'_{-i}; v) \geq u'_i(b', \beta'_{-i}; v) - \varepsilon \\ \implies & H_i(b, \beta'_{-i}) \cdot (v - b) \geq H_i(b', \beta'_{-i}) \cdot (v - b') - \varepsilon \\ \implies & H_i(b, \beta'_{-i}) \cdot (v_i - b) \geq H_i(b', \beta'_{-i}) \cdot (v_i - b') - \varepsilon + (H_i(b', \beta'_{-i}) - H_i(b, \beta'_{-i})) \cdot (v - v_i) \end{aligned}$$




Figure 1: Discrete \rightarrow Continuous




Figure 2: Continuous \rightarrow Discrete

$$\begin{aligned} \implies & u'_i(b, \beta'_{-i}; v_i) \geq u'_i(b', \beta'_{-i}; v_i) - \varepsilon - \delta \\ \implies & u_i(b, \beta_{-i}; v_i) \geq u_i(b', \beta_{-i}; v_i) - \varepsilon - \delta \end{aligned}$$

where we used (12), $|v - v_i| \leq \delta$, and the terms $H_i(b, \beta'_{-i}) \in [0, 1]$ which denote the probability for bidder i to win the item when bidding b , while the other bidders follow the strategy profile β'_{-i} in the CFPA. This completes the proof. \square

Lemma 4.3. *Given $\delta \in (0, 1)$ (in unary) and a CFPA, we can construct in polynomial time a DFPA (with the same bidding space) such that for any $\varepsilon \geq 0$, we can transform in polynomial time any ε -well-supported **monotone** MBNE of the DFPA into an $(\varepsilon + \delta)$ -PBNE of the CFPA. Furthermore, this reduction maps IPV (resp. iid) auctions to IPV (resp. iid) auctions, and symmetric equilibria to symmetric equilibria.*

Proof. Let $\delta \in (0, 1)$ be given in unary. Consider a CFPA given by continuous distributions $F'_{i,j}$ with piecewise constant density functions. Denote by $S_{i,j}$ the smallest set of points in $[0, 1]$ such that $\{0, 1\} \subseteq S_{i,j}$ and the density function of $F'_{i,j}$ is constant between any two adjacent points in $S_{i,j}$. Let $S := \cup_{i,j \in N: i \neq j} S_{i,j}$, and note that $|S|$ is polynomial in the size of the CFPA instance. Now construct a set $V \subset [0, 1]$ of polynomial size such that $S \subseteq V$ and any two adjacent points in V are separated by distance at most δ . Since $\delta \in (0, 1)$ is given in unary, this is easily achieved, e.g., by taking the union of S and a δ -grid on $[0, 1]$.

We construct a DFPA with the same bidding space B and with value spaces $V_i := V \setminus \{1\}$ for all $i \in N$. For each $i \neq j$, the discrete distribution $F_{i,j}$ over V_j is constructed from $F'_{i,j}$ as follows. For any $v_j \in V_j$, let $v_j^+ \in V$ denote the next largest element in V . Note that this is well-defined since $V = V_j \cup \{1\}$. Now we define $F_{i,j}$ to be the distribution that draws $v_j \in V_j$ with probability equal to the probability that $v \in [v_j, v_j^+]$ when $v \in [0, 1]$ is drawn according to $F'_{i,j}$. For an illustration of the induced distributions, see Figure 2. This indeed defines a valid distribution on V_j , since $0 \in V_j$ and $1 \in V$, and so the intervals $[v_j, v_j^+]$, $v_j \in V_j$, form a partition of $[0, 1]$ (ignoring endpoints). Note that the additional structure of the auction (namely IPV or iid) is kept intact.

Now consider any ε -well-supported monotone MBNE β of the DFPA. We construct a corresponding strategy profile β' in the original CFPA. For any $j \in N$, $\beta'_j : [0, 1] \rightarrow B$ is a step-function constructed as follows. We describe β'_j separately in each interval $[v_j, v_j^+]$, $v_j \in V_j$. For each $b \in \text{supp}(\beta_j(v_j))$, there is a segment of height b and length $\beta_j(v_j)(b) \cdot |v_j^+ - v_j|$. The segments are ordered inside $[v_j, v_j^+]$ in order of increasing b . This yields a non-decreasing step-function inside $[v_j, v_j^+]$. Since the intervals $[v_j, v_j^+]$ form a partition of $[0, 1]$, this defines β'_j as a step-function⁸ over $[0, 1]$. Since β_j was assumed to be

⁸In the case where β'_j is assigned two different values at a common endpoint between two adjacent intervals, we can just pick the smallest one.

monotone, it follows that β'_j is a non-decreasing step-function over all of $[0, 1]$. Furthermore, since β_j is non-overbidding, we have that $\max \text{supp}(\beta_j(v_j)) \leq v_j$ and thus $\max \beta'_j([v_j, v_j^+]) \leq v_j$, which means that β'_j is also non-overbidding. Finally, note that if β is symmetric, then so is β' .

The construction of $F_{i,j}$ from $F'_{i,j}$, and the construction of β' from β have been carefully devised to ensure that for every $i \neq j$ the two following distributions over B coincide

- Pick $v_j \in V_j$ according to $F_{i,j}$, and then output $b \in B$ according to $\beta_j(v_j)$.
- Pick $v \in [0, 1]$ according to $F'_{i,j}$ and output $b = \beta'_j(v)$.

As a result, we have that

$$u_i(b, \beta_{-i}; v_i) = u'_i(b, \beta'_{-i}; v_i) \quad (13)$$

for all $i \in N$, $b \in B$, and $v_i \in V_i$, where u_i, u'_i denote the utilities in the DFPA and CFPA respectively.

It remains to prove that β' is an $(\varepsilon + \delta)$ -PBNE of the CFPA. Consider any $i \in N$ and any $v \in [0, 1]$, and let $b = \beta'_i(v)$. We wish to show that $u'_i(b, \beta'_{-i}; v) \geq u'_i(b', \beta'_{-i}; v) - \varepsilon - \delta$ for all $b' \in B$. By construction of β'_i from β_i , there exists $v_i \in V_i$ such that $v \in [v_i, v_i^+]$ and $b \in \text{supp}(\beta_i(v_i))$. Furthermore, by construction of V_i we also have $|v - v_i| \leq \delta$. Now, since β is an ε -well-supported MBNE of the DFPA, we have that for all $b' \in B$

$$\begin{aligned} & u_i(b, \beta_{-i}; v_i) \geq u_i(b', \beta_{-i}; v_i) - \varepsilon \\ \implies & u'_i(b, \beta'_{-i}; v_i) \geq u'_i(b', \beta'_{-i}; v_i) - \varepsilon \\ \implies & H_i(b, \beta'_{-i}) \cdot (v_i - b) \geq H_i(b', \beta'_{-i}) \cdot (v_i - b') - \varepsilon \\ \implies & H_i(b, \beta'_{-i}) \cdot (v - b) \geq H_i(b', \beta'_{-i}) \cdot (v - b') - \varepsilon + (H_i(b', \beta'_{-i}) - H_i(b, \beta'_{-i})) \cdot (v_i - v) \\ \implies & u'_i(b, \beta'_{-i}; v) \geq u'_i(b', \beta'_{-i}; v) - \varepsilon - \delta \end{aligned}$$

where we used (13), $|v - v_i| \leq \delta$, and the terms $H_i(b, \beta'_{-i}) \in [0, 1]$ which denote the probability for bidder i to win the item when bidding b , while the other bidders follow the strategy profile β'_{-i} in the CFPA. This completes the proof. \square

4.3 PPAD-completeness

Next, we provide the proof of [Informal Theorem 2](#), i.e., the PPAD-completeness of computing approximate MBNE of the DFPA with subjective priors. The PPAD-membership follows directly from our equivalence result and the PPAD-membership of pure PBNE of the CFPA proven by [Filos-Ratsikas et al. \[FGHLP23\]](#). Our result does not only establish that computing ε -approximate equilibria is in PPAD, but that the membership holds even for ε -well-supported MBNE ([Definition 4](#)) which are also monotone ([Definition 3](#)). The formal statement of our membership lemma is below:

Lemma 4.4. *The problem of computing a monotone ε -well-supported MBNE of the DFPA with subjective priors is in PPAD.*

Proof. First, we observe that the problem lies in the class TFNP of total search problems that have efficiently verifiable solutions [[MP91](#)]. This is because by [Lemma 2.1](#), we can compute the utility of a given strategy in polynomial time. Then, it suffices to check if a given strategy is utility-maximizing against all possible pure deviations, see also [Remark 1](#). For the PPAD-membership, we apply [Lemma 4.2](#), which reduces the problem of finding a monotone ε -well-supported MBNE of the DFPA to that of finding a (monotone) ε -approximate PBNE of the CFPA. The latter problem is in PPAD by [[FGHLP23](#), [Theorem 4.4.](#)]. \square

We now turn our attention to the PPAD-hardness of the problem. A first observation is that our discrete-continuous equivalence result can already yield *some* PPAD-hardness result for the DFPA: starting from the PPAD-hardness result of Filos-Ratsikas et al. [FGHLP23, Theorem 5.1.], one could use Lemma 4.3 to reduce from the continuous to the discrete version. However, that (coupled with Lemma 4.1) would result in the PPAD-hardness of computing *monotone* ε -approximate MBNE. Monotonicity is a desirable property for positive results (as it makes the results stronger), but we would ideally like to prove hardness for *any* type of MBNE, monotone or not. Another caveat of the equivalence approach is that the PPAD-hardness result would require instances with relatively large bidding spaces.

To avoid these shortcomings, we provide a direct PPAD-hardness result for the problem, via a reduction from the recently proposed PPAD-complete problem PURECIRCUIT [DFHM22]. Our reduction shows computational hardness for ε -approximate equilibria that need not be monotone, and for instances with very small bidding spaces, thus making the hardness result stronger. We remark that the proof of Lemma 4.5 below first establishes PPAD-hardness for ε -well-supported MBNE on instances with small interaction degree, and then invokes Lemma 4.1 to strengthen the result to the PPAD-hardness of ε -approximate MBNE. The proof is included in Appendix D.

Lemma 4.5. *There exists a constant $\varepsilon > 0$ such that computing an ε -approximate MBNE of the DFPA with subjective priors is PPAD-hard.*

5 Symmetric Equilibria in IID Auctions

Motivated by the computational hardness results in the previous sections, we now consider the natural special case of *iid priors*, and the computation of *symmetric mixed equilibria* in such auctions. Interestingly, positive results for this case were not known before, even for continuous auctions. We show that symmetric MBNE always exist in a DFPA and we provide a PTAS for their computation.

Before we get to this result, we present the following useful lemma. We will make use of the lemma for iid priors, but it holds more generally for subjective priors. The lemma shows that one can, in polynomial time, translate an ε -approximate MBNE of the DFPA in a *shrunk* bidding space (of smaller size), to an ε' -approximate equilibrium of the auction on the original bidding space. A process similar in nature was presented in [CP23, Proof of Thm 1.2, Step 1] for the case of the CFPA in the IPV model, and *ex-ante pure* equilibria.⁹ Our proof holds for subjective priors and interim mixed equilibria, and is simpler in nature.

Lemma 5.1 (Bidding Space Shrinkage Lemma). *Consider a DFPA with bidding space B and let M be a positive integer. We can construct a bidding space $B' \subseteq B$ with cardinality $|B'| \leq M$, in time polynomial in M and the size of the input such that any ε -approximate MBNE of the auction restricted to the bidding space B' is a $(\varepsilon + \frac{1}{M})$ -approximate MBNE in the original auction.*

Proof. Given a bidding space B and a positive integer M , we construct the restricted bidding space as follows:

$$\tilde{B} = \bigcup_{\ell=0}^{M-1} \min \left\{ b \in B \mid \ell \frac{1}{M} \leq b \leq (\ell + 1) \frac{1}{M} \right\},$$

where we allow a slight abuse of notation, such that the minimum over an empty set is defined to be the empty set. To better understand how this smaller bidding space can approximate the original, it is useful to define the following notion of distance between sets, which is an instantiation of the Hausdorff distance:

$$d(B, B') := \max_{b \in B} \min_{b' \in B'} |b - b'|$$

⁹See Appendix A for the definition of the notion.

By the definition of \tilde{B} , we can see that in our case this is bounded by the size of each interval, hence:

$$d(B, \tilde{B}) \leq \frac{1}{M}. \quad (14)$$

Let \mathcal{A} be the original auction (with bidding space B) and $\tilde{\mathcal{A}}$ be the one corresponding to the new bidding space \tilde{B} . Additionally, let $\tilde{\beta}$ be an ε -approximate MBNE of $\tilde{\mathcal{A}}$. First, we show that for any (mixed) strategy profile $\tilde{\beta}$ of \tilde{B} , any bidder i with a value v_i , and any bid $b \in B$, there has to exist a $\tilde{b} \in \tilde{B}$ such that:

$$\left| u_i(\tilde{b}, \tilde{\beta}_{-i}; v_i) - u_i(b, \tilde{\beta}_{-i}; v_i) \right| \leq d(B, \tilde{B}).$$

To see why this is true, we will bound the left-hand side in the following way:

$$\begin{aligned} \left| u_i(\tilde{b}, \tilde{\beta}_{-i}; v_i) - u_i(b, \tilde{\beta}_{-i}; v_i) \right| &= \left| (v_i - \tilde{b})H(\tilde{b}, \tilde{\beta}_{-i}) - (v_i - b)H(b, \tilde{\beta}_{-i}) \right| \\ &\leq \left| b - \tilde{b} \right| \max\{H(b, \tilde{\beta}_{-i}), H(\tilde{b}, \tilde{\beta}_{-i})\} \\ &\leq \left| b - \tilde{b} \right| \leq d(B, \tilde{B}) \end{aligned} \quad (15)$$

where we have used the definition of the utility function (5) and we have bounded the function H trivially by 1 (since it is a probability).

Next, assume that $\tilde{\beta}$ is an ε -approximate MBNE of $\tilde{\mathcal{A}}$. Then, by [Remark 1](#), it should be the case that, for all $i \in N$, $v_i \in V_i$, $\tilde{b} \in \tilde{B}$, the following holds:

$$u_i(\tilde{\beta}; v_i) \geq u_i(\tilde{b}, \tilde{\beta}_{-i}; v_i) - \varepsilon \quad (16)$$

Using the bounds derived in [Equations \(14\) and \(15\)](#), we can substitute into (16) to obtain:

$$u_i(\tilde{\beta}; v_i) \geq u_i(b, \tilde{\beta}_{-i}; v_i) - \varepsilon - \frac{1}{M} \quad \forall i \in N, \forall v_i \in V_i, \forall b \in B,$$

and this is precisely (due to [Remark 1](#)) the sufficient condition for $\tilde{\beta}$ to be an $(\varepsilon + \frac{1}{M})$ -approximate MBNE of \mathcal{A} . □

We now present our positive results for the DFPA with iid priors and the computation of symmetric MBNE. We start with the following theorem which establishes the existence of symmetric and monotone equilibria for the auction in this case.

Theorem 5.2. *In every DFPA with iid priors there exists a symmetric and monotone (exact) MBNE.*

Our proof will go via the variant of the auction with continuous values, again appealing to our computational equivalence result in [Section 4.2](#). In particular, we will first show, employing Kakutani's fixed point theorem [[Kak41](#)], that the CFPA always admits an (exact) monotone PBNE. By [Lemma 4.2](#) this will imply the existence of a monotone ε -approximate MBNE for the DFPA. Finally, taking the limit of the sequence of approximate equilibria as $\varepsilon \rightarrow 0$, we will obtain the existence of an exact monotone equilibrium.

We remark that, to the best of our knowledge, general existence results for symmetric equilibria and iid priors have not been proven for either the DFPA or the CFPA. There is a plethora of results for symmetric equilibria of auctions where both the values and the bids are continuous (e.g., see [[RS81](#); [CH13](#)] and [[Kri09](#), Section 2.3]), but those do not have implications for our setting. For the case of the CFPA, existence results have only been shown for special cases [[Chw89](#)]. The proof of [Theorem 5.2](#) is in [Appendix E](#).

We now proceed to present a PTAS (*Polynomial-time approximation scheme*) for computing symmetric ε -approximate MBNE for iid priors; in our setting, this means an algorithm that on input the description of the DFPA and any $\varepsilon > 0$, finds an ε -approximate monotone MBNE of the auction in time polynomial in the description of the auction and possibly exponential in $1/\varepsilon$. For any constant $\varepsilon > 0$, this is a polynomial time algorithm for the problem.

Our strategy for constructing this algorithm will be to express the equilibrium computation problem as a system of polynomial inequalities. By known results in the literature (e.g., see [GV88]), such a system can be solved to any accuracy $\delta > 0$ in time polynomial in $1/\log(\delta)$ and $(\kappa d)^\beta$, where κ is the number of polynomials, d is their maximum degree, and β is the number of variables of the system. However, if we were to represent our equilibrium computation problem as such a system “naively”, we would end up with a running time that is exponential in both $|V|$ (the common value space for all the bidders) and $|B|$. To avoid this, we exploit the monotonicity of the equilibrium to come up with a much more succinct representation of the strategies, which we refer to as *support-representation*. This allows us to remove the exponential dependence on $|V|$, but $|B|$ still appears in the exponent of the running time. For that, we invoke Lemma 5.1 to obtain a constant-sized bidding space, at the expense of some accuracy in the efficiency of our computed equilibrium.

Theorem 5.3. *In every DFPA with m bids and n bidders having iid priors over k values, a symmetric ε -approximate mixed equilibrium can be computed in time polynomial in $\log(1/\varepsilon)$ and $(nkm)^{m^2}$.*

Proof. Fix a DFPA with $N = [n]$ bidders, value space $V_i = V = \{v_1, v_2, \dots, v_k\}$ (indexed by $j \in [k]$) for all bidders $i \in N$, bidding space $B = \{b_1, b_2, \dots, b_m\}$ (indexed by $\ell \in [m]$), and iid priors with probability mass function $\{f(v_i)\}_{i \in [k]}$.

Support-representation of monotone strategies. Let $\beta_i : V \rightarrow \Delta(B)$ be a symmetric (mixed) strategy of some bidder i . Then, β_i can be represented in the following, more succinct way, which we will call *support-representation*.

We maintain a sequence of indices over the bidding set that keeps track of which values have an “active” support on the given bid. Formally, we define a *monotone* sequence $\xi : [m] \rightarrow [k]$, such that $\xi(m) = k$, and helper functions

$$\begin{aligned} J^{\text{mixed}} &:= \xi([m]) \\ J^{\text{pure}} &:= [k] \setminus J^{\text{mixed}} \\ \ell^-(j) &:= \min \{ \ell \in [m] \mid \xi(j) = \ell \}, & \text{for } j \in J^{\text{mixed}}, \\ \ell^+(j) &:= \min \{ \max \{ \ell \in [m] \mid \xi(j) = \ell \} + 1, m \}, & \text{for } j \in J^{\text{mixed}}, \\ \ell^-(j) = \ell^+(j) = \ell(j) &:= \min \{ \ell \in [m] \mid \xi(j) > \ell \}, & \text{for } j \in J^{\text{pure}}, \\ L(j) &:= [\ell^-(j), \ell^+(j)] \cap \mathbb{N}, & \text{for } j \in [k]. \end{aligned}$$

The semantics are that strategy β_i maps value v_j to a distribution supported within $\{b_\ell\}_{\ell \in L(j)}$. In particular, under β_i , bidder i bids deterministically $b_{\ell(j)}$ when having a value in $\{v_j\}_{j \in J^{\text{pure}}}$.

Example 1. As an example of our representation, consider the sequence $\xi = (1, 1, 1, 4, 5, 5, 9)$ for a DFPA auction with 7 bids $B = \{b_1, \dots, b_7\}$ and 9 values $V = \{v_1, \dots, v_9\}$. Then, ξ corresponds to a monotone bidding rule $\beta : V \rightarrow \Delta(B)$ such that:

- $\beta(v_1)$ is supported within $\{b_1, b_2, b_3, b_4\}$
- $\beta(v_2) = \beta(v_3) = b_4$ deterministically
- $\beta(v_4)$ is supported within $\{b_4, b_5\}$
- $\beta(v_5)$ is supported within $\{b_5, b_6, b_7\}$

- $\beta(v_6) = \beta(v_7) = \beta(v_8) = \beta(v_9) = b_7$ deterministically.

To complement ξ , which is the “structural” component of our representation, we also maintain a corresponding vector of probability variables $\{p_{j,\ell}\}_{j \in J^{\text{mixed}}, \ell \in L(j)}$, with the semantics that $p_{j,\ell}$ is the probability that value v_j assigns to bid b_ℓ (under β_j). Notice that, for any given ξ , these variables are at most $2|B| = 2k$ since $|J^{\text{mixed}}| = |\xi([m])| \leq k$ and $L(j) \cap L(j') \leq 1$ for any $j, j' \in J^{\text{mixed}}$ with $j \neq j'$.

Finally, observe that if we have the aforementioned support-representation of a monotone strategy, we can easily extend it (in polynomial time) to the canonical representation of mixed strategies by defining the remaining probabilities (which are 0/1 constants), in the following way

$$p_{j,\ell} = \begin{cases} 0, & \text{if } j \in J^{\text{mixed}} \wedge \ell \neq L(j), \\ 1, & \text{if } j \in J^{\text{pure}} \wedge \ell = \ell(j), \\ 0, & \text{if } j \in J^{\text{pure}} \wedge \ell \neq \ell(j). \end{cases} \quad (17)$$

We now continue with describing our algorithm for computing an approximate symmetric equilibrium of our auction.

Guessing the support-representation structure. Let $\beta = (\beta, \beta, \dots, \beta)$ be an *exact* MBNE of our auction, where all bidders employ the same monotone¹⁰ (mixed) strategy β . Recall that such an object is guaranteed to exist, due to [Theorem 5.2](#).

The bidder strategies β of such an equilibrium will have the same support-representation ξ . Since ξ is a m -length sequence whose elements belong in $[k]$, we can exhaustively go over all such possible structural vectors that our equilibria can have in time $O(k^m)$.

The system of polynomial inequalities. Once such a structure vector ξ for the individual bidder strategies is fixed, the equilibrium β is fully determined by the probabilities $\{p_{j,\ell}\}$ of the support-representation. More precisely, β is an equilibrium if and only if the following constraints are satisfied:

- (Feasibility of strategies) For each $v_j \in V$, probabilities $\{p_{j,\ell}\}$ should give rise to a valid distribution over bids b_ℓ , i.e.

$$p_{j,\ell} \geq 0 \quad \forall j \in [k], \forall \ell \in [m] \quad (18)$$

$$\sum_{\ell=1}^m p_{j,\ell} = 1 \quad \forall j \in [k] \quad (19)$$

- (No overbidding) No bidder should bid higher than her true value with strictly positive probability, under strategy β :

$$p_{j,\ell}(v_j - b_\ell) \geq 0 \quad \forall j \in [k], \forall \ell \in [m]. \quad (20)$$

- (No improvement) For any possible true value, no bidder should have an incentive to unilaterally deviate from the (possibly randomized) bidding strategy dictated by β , to any alternative deterministic bid (see [Definition 1](#) for $\varepsilon = 0$). As a matter of fact, since β is symmetric, it is enough to consider condition (4) from the perspective of a single (arbitrary) bidder i . So, letting $u(\ell, j) := u_i(b_\ell, \beta_{-i}; v_j)$, the equilibrium condition can be written as

$$\sum_{\ell'=1}^m p_{j,\ell'} u(\ell', j) \geq u(\ell, j) \quad \forall j \in [k], \forall \ell \in [m], \quad (21)$$

¹⁰See [Definition 3](#).

where, by using Equations (5) to (8) and the fact that the bidders are iid, we can express

$$u(\ell, j) = (v_j - b_\ell) \sum_{r=0}^{n-1} \frac{1}{r+1} \binom{n-1}{r} g_\ell^r G_\ell^{n-r-1}, \quad (22)$$

with

$$g_\ell = \sum_{j=1}^k p_{j,\ell} f(v_j) \quad \forall \ell \in [m] \quad (23)$$

denoting the probability that any given bidder $i' \neq i$ bids exactly b_ℓ (see (6)) under profile β_{-i} , and

$$G_\ell = \sum_{\ell'=1}^{\ell-1} g_{\ell'} \quad \forall \ell \in [m] \quad (24)$$

the probability that she bids strictly less than b_ℓ (see (7)). Observe that in (22) the utility $u(\ell, j)$ is expressed as a polynomial, of degree at most $n - 1$, over the variables $\{g_\ell, G_\ell\}_{\ell \in [m]}$.

Let (P) denote the feasibility program, over variables $\{p_{j,\ell}, g_\ell, G_\ell\}$, defined by putting together constraints (18), (19), (20), (21) via replacing the utilities by their expression in (22), (23), and (24). Then, any solution to (P) gives rise, via the support-representation structure ξ , to a symmetric, monotone, exact MBNE of our auction.

Now observe that (P) is a system of $O(km)$ polynomial inequalities, of maximum degree n , over $O(m)$ variables; for the latter, recall that only (at most) $2m$ of the probability variables $p_{j,\ell}$ are “real” variables in (P), namely $\{p_{j,\ell}\}_{j \in J^{\text{mixed}}, \ell \in L(j)}$, since all others are constants given by (17). Therefore, an approximate solution to (P) can be found in time polynomial in n, k and exponential in m . More precisely, let $\{\tilde{p}_{j,\ell}, \tilde{g}_\ell, \tilde{G}_\ell\}$ be a set of values for the variables of (P) that are δ -near, with respect to the maximum norm, to a feasible solution $\{p_{j,\ell}^*, g_\ell^*, G_\ell^*\}$ of (P). Such a solution can be found in time polynomial in $\log(1/\delta)$ and $(nkm)^{m^2}$ – using, e.g., the results of Grigor’ev and Vorobjov [GV88, Remark, p. 38].

Rounding back the solution. We will now show how we can transform, in polynomial time, the probabilities $\tilde{\mathbf{p}} = \{\tilde{p}_{j,\ell}\}$ that we derived by (approximately) solving system (P), to another set of probabilities, let’s denote them by $\hat{\mathbf{p}} = \{\hat{p}_{j,\ell}\}$, so that they give rise to an *approximate* (mixed) equilibrium of our auction. That is, $\hat{\mathbf{p}}$ should satisfy constraints (18), (19) and (20), but we will allow (21) to be relaxed to

$$\sum_{\ell'=1}^m \hat{p}_{j,\ell'} u(\ell', j) \geq u(\ell, j) - \varepsilon. \quad (25)$$

To that end, we define

$$\hat{p}_{j,\ell} := \frac{T(\tilde{p}_{j,\ell})}{\sum_{\ell'=1}^m T(\tilde{p}_{j,\ell'})}, \quad (26)$$

where operator T is the following:

$$T(x) = \begin{cases} 0, & \text{if } x \leq \delta, \\ 1, & \text{if } x \geq 1, \\ x, & \text{otherwise.} \end{cases} \quad (27)$$

From the fact that $T(x) \geq 0$ and the normalization in (26), it is immediately apparent that our new probabilities $\hat{\mathbf{p}}$ indeed satisfy constraints (18) and (19).

Additionally, consider how the probabilities $\mathbf{p}^* = \{p_{j,\ell}^*\}$ of the *exact* solution of (P) behave, depending on the value of $(v_j - b_\ell)$ in (20). If $v_j - b_\ell \geq 0$ then, by the fact that all elements of $\hat{\mathbf{p}}$ are non-negative, (20) will be satisfied. Otherwise, if $v_j - b_\ell < 0$ it should be the case that $p_{j,\ell}^* = 0$, and therefore it must be that $\tilde{p}_{j,\ell} \leq \delta$, since $\|\mathbf{p}^* - \tilde{\mathbf{p}}\|_\infty \leq \delta$. Therefore, due to (27) it must be that $\tilde{p}_{j,\ell} = 0$, and (20) is satisfied in this case as well.

Finally, we need to establish (25). For a value $v_j \in V$ and bid $b_\ell \in B$ we will denote

$$u_{j,\ell}(\mathbf{p}) = (v_j - b_\ell) \sum_{r=0}^{n-1} \frac{1}{r+1} \binom{n-1}{r} [g_\ell(\mathbf{p})]^r [G_\ell(\mathbf{p})]^{n-r-1}, \quad (28)$$

where

$$g_\ell(\mathbf{p}) := \sum_{j=1}^k p_{j,\ell} f(v_j), \quad \text{for } \ell \in [m] \quad \text{and} \quad G_\ell(\mathbf{p}) := \sum_{\ell'=1}^{\ell-1} g_{\ell'}(\mathbf{p}).$$

We will show that, for every $\varepsilon > 0$, we can choose a $\delta = \Theta\left(\frac{\varepsilon}{mn}\right)$ such that

$$|u_{j,\ell}(\hat{\mathbf{p}}) - u_{j,\ell}(\mathbf{p}^*)| \leq \frac{\varepsilon}{2} \quad \forall j \in [k], \forall \ell \in [m]. \quad (29)$$

Note that (29) is enough in order to establish the desired validity of (25), since $u(\mathbf{p})$ is simply a rewriting of the utility $u(\ell, j)$ (see (22)) directly as a function of the probability variables $\mathbf{p} = \{p_{j',\ell'}\}$, and \mathbf{p}^* is a solution of (P) and thus satisfies (21).

We first bound the distance between $\hat{\mathbf{p}}$ and \mathbf{p}^* in the following lemma:

Lemma 5.4. *If $\delta \leq \frac{1}{3m}$, then $\|\hat{\mathbf{p}} - \mathbf{p}^*\|_\infty \leq 4\delta$.*

Proof. Since $p_{j,\ell}^* \in [0, 1]$ for all $j \in [k]$, $\ell \in [m]$, and $\|\tilde{\mathbf{p}} - \mathbf{p}^*\|_\infty \leq \delta$, we can deduce that

$$-\delta \leq \tilde{p}_{j,\ell} \leq 1 + \delta \quad \forall j \in [k], \forall \ell \in [m].$$

Therefore, by the definition of operator T in (27) we get that

$$|T(\tilde{p}_{j,\ell}) - \tilde{p}_{j,\ell}| \leq \delta \quad \forall j \in [k], \forall \ell \in [m]. \quad (30)$$

Therefore

$$\begin{aligned} \left| \sum_{\ell=1}^m T(\tilde{p}_{j,\ell}) - 1 \right| &= \left| \sum_{\ell=1}^m T(\tilde{p}_{j,\ell}) - \sum_{\ell=1}^m p_{j,\ell}^* \right| \\ &\leq \sum_{\ell=1}^m |T(\tilde{p}_{j,\ell}) - \tilde{p}_{j,\ell}| \\ &\leq \sum_{\ell=1}^m |T(\tilde{p}_{j,\ell}) - \tilde{p}_{j,\ell}| + \sum_{\ell=1}^m |p_{j,\ell}^* - \tilde{p}_{j,\ell}| \\ &\leq m \cdot \delta + m \cdot \delta = 2m\delta \end{aligned} \quad (31)$$

Now, for all $j \in [k]$ and $\ell \in [m]$ we can bound:

$$|\hat{p}_{j,\ell} - \tilde{p}_{j,\ell}| \stackrel{(26)}{=} \left| \frac{T(\tilde{p}_{j,\ell})}{\sum_{\ell'=1}^m T(\tilde{p}_{j,\ell'})} - \tilde{p}_{j,\ell} \right| \quad (32)$$

$$\leq \max \left\{ \frac{1}{\left| \sum_{\ell'=1}^m T(\tilde{p}_{j,\ell'}) \right|}, 1 \right\} |T(\tilde{p}_{j,\ell}) - \tilde{p}_{j,\ell}| \quad (33)$$

$$\stackrel{(30),(31)}{\leq} \max \left\{ \frac{1}{|1-2m\delta|}, 1 \right\} \delta \quad (34)$$

$$\leq 3\delta, \quad (35)$$

where for the last inequality we have used the fact that $0 < \delta \leq \frac{1}{3m}$, and for the first one the fact that $|\lambda x - \mu y| \leq \max\{|\lambda|, |\mu|\} |x - y|$ for all reals λ, μ, x, y .

Finally, we can derive that:

$$\|\hat{\mathbf{p}} - \mathbf{p}^*\|_\infty \leq \|\tilde{\mathbf{p}} - \hat{\mathbf{p}}\|_\infty + \|\tilde{\mathbf{p}} - \mathbf{p}^*\|_\infty \stackrel{(35)}{\leq} 3\delta + \delta = 4\delta.$$

□

We are now ready to bound the distance of the utilities under the two profiles $\hat{\mathbf{p}}$ and \mathbf{p}^* in (29). Using that $\|\hat{\mathbf{p}} - \mathbf{p}^*\|_\infty \leq 4\delta$ from Lemma 5.4, we can now provide the bounds:

$$|g(\hat{\mathbf{p}}) - g(\mathbf{p}^*)| = \left| \sum_{j'=1}^k \hat{p}_{j',\ell} f(v_{j'}) - \sum_{j'=1}^k p_{j',\ell}^* f(v_{j'}) \right| \leq \sum_{j'=1}^k \left| \hat{p}_{j',\ell} - p_{j',\ell}^* \right| f(v_{j'}) \leq 4\delta, \quad (36)$$

since $\sum_{j'=1}^k f(v_{j'}) = 1$, and so

$$|G(\hat{\mathbf{p}}) - G(\mathbf{p}^*)| \leq \sum_{\ell'=1}^{\ell-1} |g(\hat{\mathbf{p}}) - g(\mathbf{p}^*)| \leq 4\delta |\ell - 1| \leq 4\delta m. \quad (37)$$

Also, note that for $\mathbf{p} \in \{\hat{\mathbf{p}}, \mathbf{p}^*\}$, it is

$$G_\ell(\mathbf{p}) = \sum_{\ell'=1}^{\ell-1} g_{\ell'}(\mathbf{p}) \leq \sum_{\ell'=1}^m \sum_{j=1}^k p_{j,\ell'} f(v_j) = \sum_{j=1}^k \left(\sum_{\ell'=1}^m p_{j,\ell'} \right) f(v_j) = \sum_{j=1}^k f(v_j) = 1,$$

since both $\hat{\mathbf{p}}$ and \mathbf{p}^* satisfy (19).

Next, we rewrite the utilities in (28) as

$$\begin{aligned} u_{j,\ell}(\mathbf{p}) &= \frac{v_j - b_\ell}{n \cdot g_\ell(\mathbf{p})} \sum_{r=0}^{n-1} \binom{n}{r+1} [g_\ell(\mathbf{p})]^{r+1} [G_\ell(\mathbf{p})]^{n-r-1} \\ &= \frac{v_j - b_\ell}{n \cdot g_\ell(\mathbf{p})} \sum_{r=1}^n \binom{n}{r} [g_\ell(\mathbf{p})]^r [G_\ell(\mathbf{p})]^{n-r} \\ &= \frac{v_j - b_\ell}{n \cdot g_\ell(\mathbf{p})} [(G_\ell(\mathbf{p}) + g_\ell(\mathbf{p}))^n - G_\ell^n(\mathbf{p})] \\ &= \frac{v_j - b_\ell}{n \cdot g_\ell(\mathbf{p})} [G_{\ell+1}^n(\mathbf{p}) - G_\ell^n(\mathbf{p})] \\ &= \frac{v_j - b_\ell}{n} \sum_{r=0}^{n-1} G_{\ell+1}^r(\mathbf{p}) G_\ell^{n-1-r}(\mathbf{p}). \end{aligned}$$

Using this, we can finally bound the distance of the utilities by

$$\begin{aligned} |u_{j,\ell}(\hat{\mathbf{p}}) - u_{j,\ell}(\mathbf{p}^*)| &\leq \frac{1}{n} \sum_{r=0}^{n-1} |G_{\ell+1}^r(\hat{\mathbf{p}}) G_\ell^{n-1-r}(\hat{\mathbf{p}}) - G_{\ell+1}^r(\mathbf{p}^*) G_\ell^{n-1-r}(\mathbf{p}^*)| \\ &\leq \max_{r=0,1,\dots,n-1} |G_{\ell+1}^r(\hat{\mathbf{p}}) G_\ell^{n-1-r}(\hat{\mathbf{p}}) - G_{\ell+1}^r(\mathbf{p}^*) G_\ell^{n-1-r}(\mathbf{p}^*)| \end{aligned}$$

$$\begin{aligned}
&\leq \max_{r=0,1,\dots,n-1} |G_\ell^r(\hat{\mathbf{p}}) - G_\ell^r(\mathbf{p}^*)| \\
&\leq \max_{r=0,1,\dots,n-1} r |G_\ell(\hat{\mathbf{p}}) - G_\ell(\mathbf{p}^*)| \\
&\leq 4\delta mn,
\end{aligned}$$

where the third and fourth inequalities hold due to the fact that $G_\ell(\hat{\mathbf{p}}), G_\ell(\mathbf{p}^*) \in [0, 1]$ for all $\ell \in [m]$, and the last one due to (37). Taking $\delta = \frac{\varepsilon}{8mn}$ we can indeed satisfy (29). \square

We can now state our PTAS result.

Theorem 5.5. *In every DFPA with iid priors, and any constant $\varepsilon > 0$, a symmetric and monotone ε -approximate mixed equilibrium can be computed in polynomial time.*

Proof. Fix an iid DFPA auction \mathcal{A} with n bidders, k values, and bidding space B . Let $\varepsilon > 0$ be an arbitrary constant. Choose integer $M := \lceil 2/\varepsilon \rceil$, so that $\frac{1}{M} \leq \frac{\varepsilon}{2}$. Next, we construct a subspace $B' \subseteq B$ of our original bidding space B , with constant cardinality $|B'| = M$, as dictated in the proof of the shrinkage Lemma 5.1. Notice that B' can be constructed in polynomial time (in the description of \mathcal{A}).

Now, let \mathcal{A}' denote the DFPA that results from \mathcal{A} if we simply replace bidding space B of \mathcal{A} by B' (everything else, namely bidders, value space and belief distributions, remaining the same). Obviously, \mathcal{A}' is still an iid auction and thus, if we define $\varepsilon' := \varepsilon/2$, we can deploy the algorithm presented in Theorem 5.3 in order to compute a symmetric and monotone ε' -approximate MBNE of \mathcal{A}' , let's denote it by β^* , in time polynomial in $\log(1/\varepsilon') = O(\log(1/\varepsilon))$ and in $(nkM)^M = \text{poly}(n, k)$, since M is a constant. Due to Lemma 5.1, again, we know that β^* is guaranteed to be an $(\varepsilon' + \frac{1}{M})$ -approximate MBNE of the original auction \mathcal{A} . Our proof is concluded by noting that $\varepsilon' + \frac{1}{M} \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon$. \square

6 Correlated Equilibria

Another well-studied solution concept in Game Theory is that of Correlated Equilibria (CE), first introduced by Aumann in his seminal paper [Aum74]. Intuitively, these capture the idea of each player choosing her action after being given a recommendation from a “mediator”, ensuring that at a CE no player has an incentive to deviate from their recommended strategy. This is a broader solution concept which in several cases has proven to be amenable to polynomial-time algorithms, e.g., see [PR08; JL15]. Given our PPAD-completeness result for computing MBNE in Section 4.3, it is natural to explore CE of the DFPA, in a quest for potential tractability results.

Before we do that however, we have to come up with an appropriate definition: CE are typically defined in complete information, normal form games, whereas the DFPA is a Bayesian game of incomplete information. The appropriate extension of the notion of CE to Bayesian games has been discussed extensively in the literature [Aum87; For93; For06; For23; BM13; BM16], and several natural notions have been proposed. In this paper, we will adopt the standard notion that defines a CE of the original DFPA by means of the induced *type-agent* representation, adopted also by works in computer science, e.g., see [HST15; AB24]. This representation is a generic way of interpreting any Bayesian game as a traditional, strategic-form game; see, [Mye97, pp. 73–74, 127–128] and [JR01, Definition 7.11]. Historically, the representation has its roots directly in the seminal paper of [Har67, p. 1814], where it is referred to as the *Selten* or *posterior-lottery* model. Formally, the type-agent representation is defined as follows:

Definition 6 (Type-agent representation of the DFPA). Fix a DFPA \mathcal{A} with: bidders $N = \{1, 2, \dots, n\}$; bidding set $B = \{b_1, b_2, \dots, b_m\}$; value sets $V_i = \{v_{i,1}, v_{i,2}, \dots, v_{i,k_i}\}$, for $i \in N$; and value distributions $F_{i,j}$ with probability density functions

$$\{f_{i,j}(v_{j,1}), f_{i,j}(v_{j,2}), \dots, f_{i,j}(v_{j,k_j})\}, \quad \text{for } i, j \in N, i \neq j,$$

given explicitly in the input.

Let \mathcal{G} be the normal form game with: players $\tilde{N} = \{(i, v) \mid i \in N, v \in V_i\}$; for all players $(i, v) \in \tilde{N}$, the same set of (pure) strategies $S_{(i,v)} = B$; and, for all outcomes $\mathbf{s} \in B^{\tilde{N}}$ and players $(i, v) \in \tilde{N}$, a payoff

$$\pi_{(i,v)}(\mathbf{s}) = u_i(\mathbf{s}(i, v), \hat{\boldsymbol{\beta}}_{-i}; v), \quad (38)$$

where u_i 's are the standard DFPA utilities (as given in (2)) and $\hat{\boldsymbol{\beta}}_{-i}$ are *pure* bidding strategies of \mathcal{A} given by

$$\hat{\beta}_j(v_j) = \mathbf{s}(j, v_j), \quad (39)$$

for all $j \in N \setminus i$ and $v_j \in V_j$.

Notice that (39) gives rise to a natural one-to-one correspondence between pure (and, by extension, mixed as well) equilibria of auction \mathcal{A} and the normal-form game \mathcal{G} . Furthermore, from (38) it is apparent that this correspondence preserves the utility/payoffs between games. In particular, this means that there is a one-to-one correspondence between ε -MBNE of \mathcal{A} and ε -approximate (additive) mixed Nash equilibria of \mathcal{G} . In fact, this has been the common approach for showing existence of equilibria of Bayesian games in general, by applying Nash's theorem to the type-agent game, which is indeed in normal form.

Given Definition 6, we can now define the correlated equilibria of the DFPA to be precisely the correlated equilibria of the corresponding type-agent game.

Definition 7 (Correlated Equilibrium of a DFPA). Given a DFPA \mathcal{A} and the corresponding type-agent representation game \mathcal{G} , a correlated equilibrium is a distribution σ on the set of outcomes $B^{\tilde{N}}$ such that for every player $\tilde{n} \in \tilde{N}$, strategy $s_{\tilde{n}} \in S_{\tilde{n}}$, and every deviation $s'_{\tilde{n}} \in S_{\tilde{n}}$,

$$\mathbb{E}_{s \sim \sigma} [\pi_{\tilde{n}}(\mathbf{s}) \mid s_{\tilde{n}}] \geq \mathbb{E}_{s \sim \sigma} [\pi_{\tilde{n}}(s'_{\tilde{n}}, \mathbf{s}_{-\tilde{n}}) \mid s_{\tilde{n}}]$$

One of the most celebrated results on correlated equilibria is due to [PR08], as well as follow-up work in [JL15], and it states the following:

Theorem 6.1 ([PR08]). *In any succinct game of polynomial type for which the players' utilities can be computed in polynomial time, one can compute a correlated equilibrium in polynomial time.*

Informally, succinct games are games that can be represented without explicitly listing the payoffs of the players on all possible outcomes (that, in general, may be exponentially many), but rather via some alternative, efficient way. We refer the reader to [PR08] for a formal definition of the class.

Having established these important notions and results, the main theorem of this section follows rather straightforwardly, by applying Lemma 2.1.

Theorem 6.2. *A correlated equilibrium of a DFPA can be computed in polynomial time.*

Proof. Given Theorem 6.1, it suffices to show that (1) the induced game \mathcal{G} is succinct and (2) we can compute the (expected) utilities of the players in polynomial time. First, observe that the number of players and (pure) strategies of \mathcal{G} are polynomial in the description size $|\mathcal{A}|$ of the original DFPA and, furthermore, due to (38) and Lemma 2.1, the payoff functions $\pi_{(i,v)}$ are computable in polynomial (with respect to $|\mathcal{A}|$) time. Therefore, the class of games \mathcal{G} that could arise through our reduction are succinctly representable (of polynomial type). Additionally, Lemma 2.1 guarantees that the (expected) payoffs of \mathcal{G} on mixed strategy profiles are also polynomial-time computable, hence, the second condition is also satisfied. \square

7 Conclusion and Future Work

In this paper, we considered the computational complexity of equilibrium computation in first-price auctions with discrete values and discrete bids. We established computational completeness results for deciding the existence of pure equilibria and for finding mixed equilibria in general. We also provided positive results for the natural case of iid bidders and symmetric equilibria, as well as for correlated equilibria. At the heart of several of our results lies a novel equivalence between mixed equilibria of the discrete auction and pure equilibria of the continuous auction. We believe that this result may prove to be quite useful in the future, as it allows one to choose which of the two settings to focus on, and translate results automatically to the other.

Perhaps the most interesting, and seemingly rather challenging avenue for future work is to study the computational complexity of equilibrium computation in the discrete auction in the IPV setting. Another interesting question is to explore alternative notions of correlated equilibria defined in the literature for Bayesian games (e.g., see [For93; For06; For23]) and explore the computational relation between them in the context of first-price auctions. Finally, one could explore mixed equilibria for auctions with continuous values, after coming up with an appropriate representation of the bidders' strategies in this setting.

Acknowledgements

We are very grateful to Diogo Poças for useful ideas that served as a starting point for the NP-hardness proof.

Appendix

A A Brief Discussion on Equilibrium Notions

In the literature of the computational complexity of strategic games, the term “ ε -well-supported equilibrium” has been used to refer to ε -approximate mixed equilibria in which every pure strategy in the support of the equilibrium strategy is utility-maximising, exactly like in [Definition 4](#). Recently, Chen and Peng [[CP23](#)] ambiguously used the term “ ε -well-supported” to refer to *pure interim* Bayes-Nash equilibria. An interim equilibrium is one as defined in our [Definition 1](#), in which a player chooses her strategy given her value v_i , which is known to her. In the terminology of auction theory, one could also define an alternative equilibrium notion called *ex-ante* Bayes-Nash equilibrium. In such an equilibrium, each bidder also has some uncertainty about her own value, which she draws from a distribution $F_{i,i}$. We provide the formal definition below.

Definition 8 (ε -approximate *ex-ante* mixed Bayes-Nash equilibrium of the DFPA). Let $\varepsilon \geq 0$. A (mixed) strategy profile $\boldsymbol{\beta} = (\beta_1, \dots, \beta_n)$ is an (interim) ε -approximate mixed Bayes-Nash equilibrium (MBNE) of the DFPA if for any bidder $i \in N$,

$$\mathbb{E}_{v_i \sim F_{i,i}} [u_i(\beta_i(v_i), \boldsymbol{\beta}_{-i}; v_i)] \geq \mathbb{E}_{v_i \sim F_{i,i}} [u_i(\boldsymbol{\gamma}, \boldsymbol{\beta}_{-i}; v_i)] - \varepsilon \quad \text{for all } \boldsymbol{\gamma} \in \Delta(B).$$

By [Definition 4](#) and [Definition 1](#), it follows that an ε -well-supported equilibrium is also an ε -approximate Bayes-Nash equilibrium. Additionally, from [Definition 8](#) above and [Definition 1](#), it follows that an interim Bayes-Nash equilibrium is also an *ex-ante* Bayes-Nash equilibrium. This implies that positive results are stronger for interim equilibria. In particular, the PTAS of Chen and Peng [[CP23](#)] for the IPV setting, which is stated for *ex-ante* equilibria, does not translate to our setting via our computational equivalence result.

We remark that we find the notion of interim equilibria more natural, and more in line with the standard definitions (e.g., see [[Kri09](#); [Mye97](#)]), and advocate the use of terminology which is consistent with the auction literature, e.g., see [[MR85](#); [MR00](#); [Leb96](#); [Ath01](#)].

B Non-Existence of Approximate Pure Equilibria

In this section we prove [Theorem 3.1](#). We start by establishing a technical lemma that would be useful in our main proof:

Lemma B.1. *For any integer $M \geq 10$,*

$$\frac{q}{2}M \leq q(M-2) \leq \frac{q+1}{2}(M-3) \leq M-4 \leq q(M-1) \leq \frac{q+1}{2}(M-2) \leq M-3 \leq \frac{q+1}{2}(M-1),$$

where

$$q = \frac{M-3}{M-1} - \frac{2}{3M}.$$

Proof. First, we show that, for $M \geq 10$, the following holds:

$$1 - \frac{3}{M-1} \leq q \leq 1 - \frac{2}{M-3}. \tag{40}$$

Indeed, we can see that:

$$1 - \frac{3}{M-1} \leq q \iff 1 - \frac{3}{M-1} \leq \frac{M-3}{M-1} - \frac{2}{3M} \iff M > 1$$

and

$$q \leq 1 - \frac{2}{M-3} \iff \frac{M-3}{M-1} - \frac{2}{3M} \leq 1 - \frac{2}{M-3} \iff M^2 - 10M + 3 \geq 0,$$

which holds for $M \geq \sqrt{22} + 5 \approx 9.690$.

To prove the main chain of inequalities in the statement of our lemma, again by performing basic algebraic operations, we get:

$$\frac{q}{2}M \leq q(M-2) \iff M \geq 4;$$

$$q(M-2) \leq \frac{q+1}{2}(M-3) \iff q \leq \frac{M-3}{M-1}$$

which holds since $q = \frac{M-3}{M-1} - \frac{2}{3M}$;

$$\frac{q+1}{2}(M-3) \leq M-4 \iff q \leq 1 - \frac{2}{M-3}$$

which holds due to (40);

$$M-4 \leq q(M-1) \iff q \geq 1 - \frac{3}{M-1}$$

which holds due to (40);

$$\leq q(M-1) \leq \frac{q+1}{2}(M-2) \iff q \leq 1 - \frac{2}{M}$$

which holds since, from (40), we have that $q \leq 1 - \frac{2}{M-3} < 1 - \frac{2}{M}$;

$$\frac{q+1}{2}(M-2) \leq M-3 \iff q \leq 1 - \frac{2}{M-2}$$

which holds since, from (40), we have that $q \leq 1 - \frac{2}{M-3} < 1 - \frac{2}{M-2}$;

$$M-3 \leq \frac{q+1}{2}(M-1) \iff q \geq 1 - \frac{4}{M-1}$$

which holds since, from (40), we have that $q \leq 1 - \frac{3}{M-1} > 1 - \frac{4}{M-1}$. □

We are now ready to proceed with the proof of our theorem:

Proof of Theorem 3.1. Let $M := 12$ and fix any ε such that

$$0 \leq \varepsilon < \frac{1}{3M} - \frac{2}{M^2} = \frac{1}{72}. \quad (41)$$

Consider a DFPA with two bidders with value spaces $V_1 = V_2 = \{0, 1\}$, bidding space $B = \{0, \frac{1}{M}, \frac{2}{M}, \dots, 1\}$, and identical prior distributions with probability mass function

$$f(v) = \begin{cases} \frac{M-3}{M-1} - \frac{2}{3M} =: q, & \text{if } v = 0, \\ 1 - q, & \text{if } v = 1. \end{cases} \quad (42)$$

Notice that this distribution is well-defined, since $q \in (0, 1)$ for $M \geq 10$ (see (40) in the proof of Lemma B.1).

First, we will show that no profile can be an ε -approximate equilibrium, unless $\beta_1(0) = \beta_2(0) = 0$. Assume (without loss) that $\beta_1(0) \geq \beta_2(0)$, and let $\beta_1(0) = b \geq \frac{1}{M}$. We will prove that bidder 1 can get an ε -improvement by deviating to $\beta_1(0) = 0$. Indeed, the probability that bidder 1 gets the item by reporting b is at least

$$\frac{1}{2} \Pr [\beta_2(v_2) \leq b] \geq \frac{1}{2} \Pr [v_2 = 0] = \frac{q}{2}$$

so, keeping everything else fixed and conditioning on the true value of bidder 1 being $v_1 = 0$, the improvement in the utility is:

$$\begin{aligned}
u_1(\beta_1(0) = 0) - u_1(\beta_1(0) = b) &= 0 - \Pr[\text{bidder 1 wins given report } b] \cdot (0 - b) \\
&\geq \frac{q}{2}b \geq \frac{q}{2} \frac{1}{M} \\
&= \left[\frac{M-3}{2(M-1)} - \frac{1}{3M} \right] \frac{1}{M}, \quad \text{by substituting } q \text{ from (42),} \\
&> \left[\frac{M-3}{3M} - \frac{1}{3M} \right] \frac{1}{M} \\
&= \frac{1}{3M} - \frac{4}{3} \frac{1}{M^2} > \frac{1}{3M} - 2 \frac{1}{M^2} = \varepsilon.
\end{aligned}$$

Thus, from now on we fix $\beta_1(0) = \beta_2(0) = 0$, and we use the shorthands $\beta_1 = \beta_1(1)$, $\beta_2 = \beta_2(1)$. We will now show that, for any bidder $i = 1, 2$:

(a) if $\beta_{-i} \in \{0, \frac{1}{M}, \frac{2}{M}\}$, then

$$u_i\left(\beta_i = \beta_{-i} + \frac{1}{M}; v_i = 1\right) - u_i(\beta_i = b_i; v_i = 1) > \varepsilon \quad \text{for all } b_i \neq \beta_{-i} + \frac{1}{M}$$

(b) if $\beta_{-i} \in \{\frac{3}{M}, \frac{4}{M}, \dots, 1\}$, then

$$u_i\left(\beta_i = \frac{1}{M}; v_i = 1\right) - u_i(\beta_i = b_i; v_i = 1) > \varepsilon \quad \text{for all } b_i \neq \frac{1}{M}.$$

This is enough to establish the nonexistence of an ε -approximate equilibrium, since it shows that for any choice of $\beta_1, \beta_2 \in \{0, \frac{1}{M}, \frac{2}{M}, \dots, 1\}$ there will always exist an ε -improving deviation. To simplify the notation, we will drop the conditioning on $v_i = 1$ for the rest of the proof, and also denote

$$\Delta(z) = \Delta(z, \beta_{-i}) = u_i(\beta_i = z; v_i = 1) - \max_{b_i \neq z} u_i(\beta_i = b_i; v_i = 1),$$

so that cases (a) and (b) above can be expressed simply as

$$(a) \Delta\left(\beta_{-i} + \frac{1}{M}\right) > \varepsilon \text{ for } \beta_{-i} \in \{0, \frac{1}{M}, \frac{2}{M}\}$$

$$(b) \Delta\left(\frac{1}{M}\right) > \varepsilon \text{ for } \beta_{-i} \in \{\frac{3}{M}, \frac{4}{M}, \dots, 1\}.$$

First, observe that

$$u_i(\beta_i = b_i) = 1 - b_i \quad \text{for all } b_i \geq \beta_{-i} + \frac{1}{M}$$

Next, for the first case (a), we perform a case analysis depending on $\beta_{-i} = 0, \frac{1}{M}, \frac{2}{M}$:

– if $\beta_{-i} = 0$, then $u_i(\beta_i = 0) = \frac{1}{2}(1 - 0) = \frac{1}{2}$ and so

$$\Delta\left(\frac{1}{M}\right) = 1 - \frac{1}{M} - \max\left\{\frac{1}{2}, 1 - \frac{2}{M}\right\} = \min\left\{\frac{1}{2} - \frac{1}{M}, \frac{1}{M}\right\} = \frac{1}{M} > \varepsilon.$$

– if $\beta_{-i} = \frac{1}{M}$, then

$$u_i(\beta_i = 0) = \frac{1}{2}q(1 - 0) = q \frac{1}{2}$$

and

$$u_i\left(\beta_i = \frac{1}{M}\right) = \left[q + \frac{1}{2}(1-q)\right] \left(1 - \frac{1}{M}\right) = \frac{q+1}{2} \left(1 - \frac{1}{M}\right).$$

So

$$\begin{aligned} \Delta\left(\frac{2}{M}\right) &= 1 - \frac{2}{M} - \frac{1}{M} \max\left\{q\frac{M}{2}, \frac{q+1}{2}(M-1), M-3\right\} \\ &= 1 - \frac{2}{M} - \frac{1}{M} \frac{q+1}{2}(M-1), && \text{by Lemma B.1,} \\ &= \frac{1}{3M} - \frac{1}{3} \frac{1}{M^2}, && \text{by substituting } q \text{ from (42),} \\ &> \frac{1}{3M} - 2\frac{1}{M^2} = \varepsilon. \end{aligned}$$

– if $\beta_{-i} = \frac{2}{M}$, then

$$u_i(\beta_i = 0) = \frac{1}{2}q(1-0) = q\frac{1}{2},$$

$$u_i\left(\beta_i = \frac{1}{M}\right) = q\left(1 - \frac{1}{M}\right)$$

and

$$u_i\left(\beta_i = \frac{2}{M}\right) = \frac{q+1}{2} \left(1 - \frac{2}{M}\right).$$

So

$$\begin{aligned} \Delta\left(\frac{3}{M}\right) &= 1 - \frac{3}{M} - \frac{1}{M} \max\left\{q\frac{M}{2}, q(M-1), \frac{q+1}{2}(M-2), M-4\right\} \\ &= 1 - \frac{3}{M} - \frac{1}{M} \frac{q+1}{2}(M-2), && \text{by Lemma B.1,} \\ &= \frac{M^2 - 6M + 2}{3(M-1)M^2}, && \text{by substituting } q \text{ from (42),} \\ &> \frac{M^2 - 6M}{3M^3} = \frac{1}{3M} - \frac{2}{M^2} = \varepsilon. \end{aligned}$$

For the second case (b), it is

$$u_i(\beta_i = 0) = \frac{1}{2}q(1-0) = q\frac{1}{2};$$

$$u_i(\beta_i = b_i) = q(1 - b_i) \quad \text{for all } b_i = \frac{1}{M}, \dots, \beta_{-i} - \frac{1}{M};$$

$$u_i(\beta_i = \beta_{-i}) = \frac{q+1}{2}(1 - \beta_{-i}).$$

Therefore,

$$\begin{aligned} \Delta\left(\frac{1}{M}\right) &= q\left(1 - \frac{1}{M}\right) - \frac{1}{M} \max\left\{q\frac{M}{2}, q(M-2), \frac{q+1}{2}(M - M\beta_{-i}), M - M\beta_{-i} - 1\right\} \\ &\geq q\left(1 - \frac{1}{M}\right) - \frac{1}{M} \max\left\{q\frac{M}{2}, q(M-2), \frac{q+1}{2}(M-3), M-4\right\}, && \text{since } \beta_{-i} \geq \frac{3}{M}, \\ &= q\left(1 - \frac{1}{M}\right) - \frac{1}{M}(M-4), && \text{by Lemma B.1,} \\ &= \frac{1}{3M} + \frac{2}{3M^2} > \frac{1}{3M} - \frac{2}{M^2} = \varepsilon, \end{aligned}$$

where the penultimate equality follows by substituting q from (42). □

C NP-Hardness of Deciding Approximate Equilibria

In this section, we provide the full proof for the NP-hardness of the problem. In our reduction, we consider an instance of the problem with bidding space $B = \{0, 1, 2, 3\}$ and a common value space for all bidders $V = \{0, \frac{9}{4}, 10\}$. Notice that in this case V and B do not lie inside $[0, 1]$ as we originally defined. However, it is trivial to see, by the definition of the utility, that if we divide all elements of both sets by the maximum (10 in this case), then they are all in $[0, 1]$ as required, and, additionally, the comparisons between the corresponding utilities in our proof still hold, as they have all been multiplied by $1/10$. This means that the value of ε that we will compute for the hardness of approximation will also have to be multiplied by $1/10$ at the end of the proof.

Let a DFPA instance with $B = \{0, 1, 2, 3\}$ and a common value space for all bidders, $V = \{0, \frac{9}{4}, 10\}$. In this context, the bidding strategy of a bidder i can be efficiently represented by one of the $4^3 = 64$ functions of type $V \rightarrow B$. We construct a reduction from the Circuit Satisfiability problem (CIRCUIT-SAT), which is NP-complete [GJ79].

Definition 9. Given a Boolean circuit C , CIRCUIT-SAT is the problem of deciding whether there is an assignment to its inputs that makes its output *true*.

Our construction has 3 basic parts. Given an instance of CIRCUIT-SAT, we transform it to an instance of the problem of calculating a PBNE in a DFPA as follows:

- We map *false* and *true* to two specific strategies, which we will denote as s_0 and s_1 .
- For each logic gate in the circuit, we introduce new bidders (along with their subjective prior distributions) to the DFPA, such that the best-response of the bidder corresponding to the gate’s output given the gate’s input bidder(s) reflects the type of gate (with respect to the mapping of *false* and *true* to s_0 and s_1). This ensures that a PBNE can only exist if the gates are correctly simulated.
- We introduce two extra bidders such that the conditions of a PBNE are not satisfied between them unless the bidder corresponding to the output of the circuit represents the value *true*.

C.1 Simulating Circuit Gates

We know that the NAND gate is functionally complete, i.e. all possible boolean functions can be represented by circuits using only NAND gates. Since we can construct a NAND from an OR and a NOT gate, we can transform (in polynomial time) any circuit of an instance of CIRCUIT-SAT to one that contains only OR and NOT gates. Therefore, it suffices to reduce from these new instances, which are also hard to solve. To simulate these in our reduction, we will introduce 2 corresponding gadgets, which will be used in combination with an extra gadget we define, the *projection* gadget. The *projection* gadget has one input and one output bidder, and it is used for appropriately mapping any strategies different from s_0 and s_1 that might arise as outputs of the OR and NOT gadgets back to either s_0 or s_1 . The values of *false* and *true* will be encoded by a bidder’s strategy as follows:

- *false* is encoded by bidding strategies such that $\hat{\beta}(0) = 0$, $\hat{\beta}(\frac{9}{4}) = 1$, and $\hat{\beta}(10) = 2$;
- *true* is encoded by bidding strategies such that $\hat{\beta}(0) = 0$, $\hat{\beta}(\frac{9}{4}) = 2$, and $\hat{\beta}(10) = 3$.

To simplify notation, we denote these strategies as $s_0 = (0, 1, 2)$ for *false* and $s_1 = (0, 2, 3)$ for *true*.

C.1.1 Projection gadget

We begin by designing the *projection* gadget, as we will need it to simulate the OR and NOT gates. Suppose bidder i encodes a boolean value via her strategy $\hat{\beta}_i$. Consider a bidder j whose subjective prior is given as follows: F_{ji} is the discrete distribution that assumes value 0 with probability $6/100$ and value $9/4$ with probability $94/100$; for any other $i' \neq i, j$, $F_{ji'}$ is the discrete distribution that assumes value 0 with probability 1. We will now perform a case analysis to calculate j 's best-response to i 's strategies. Whenever i plays s_0 or s_1 , j will have to copy her strategy in order to best-respond. Additionally, all other strategies would result in j 's utility being at least $\frac{1}{18}$ worse. We also provide j 's best-response to strategies other than s_0 and s_1 , as they will be required for the proof of [Lemma C.4](#). As F_{ji} has probability mass only on 0 and $9/4$, we can ignore i 's strategy when having value 10; we will denote this by x below. To begin with, j 's best-response must satisfy $\hat{\beta}_j(0) = 0$ due to the no-overbidding assumption. In what follows, we allow a slight abuse of notation where $u_j(b; v)$ denotes the expected utility that bidder j receives (from her point of view) when having value v and choosing to bid b - this is in contrast to the usual definition where we also include the other bidders' strategies, however here we explicitly list their strategies in the description of each different case. We consider the following ones:

1. Suppose bidder i plays the strategy $(0, 0, x)$. We can compute the winning probabilities and utilities for bidder j when bidding 0, 1, 2, or 3, and when having value $9/4$ or 10:

b	0	1	2	3
$H_j(b)$	$\frac{1}{n}$	1	1	1
$u_j(b; 9/4)$	$\frac{9}{4n}$	$\frac{5}{4}$	$\frac{1}{4}$	–
$u_j(b; 10)$	$\frac{10}{n}$	9	8	7

Since $n \geq 2$, we have the bounds $u_j(0; 9/4) \leq \frac{9}{8}$, $u_j(0; 10) \leq 5$. Comparing these values, we conclude that it is optimal for bidder j to bid $\hat{\beta}_j(9/4) = 1$, $\hat{\beta}_j(10) = 1$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

2. Suppose bidder i plays some strategy of the form $(0, 1, x)$. We can compute the winning probabilities and utilities for bidder j when bidding 0, 1, 2, or 3, and when having value $9/4$ or 10:

b	0	1	2	3
$H_j(b)$	$\frac{3}{50n}$	$\frac{53}{100}$	1	1
$u_j(b; 9/4)$	$\frac{27}{200n}$	$\frac{53}{80}$	$\frac{1}{4}$	–
$u_j(b; 10)$	$\frac{3}{5n}$	$\frac{477}{80}$	8	7

Since $n \geq 2$, we have the bounds $u_j(0; 9/4) \leq \frac{27}{400}$, $u_j(0; 10) \leq \frac{3}{10}$. Comparing these values, we conclude that it is optimal for bidder j to bid $\hat{\beta}_j(9/4) = 1$, $\hat{\beta}_j(10) = 2$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

3. Suppose bidder i plays some strategy of the form $(0, 2, x)$. We can compute the winning probabilities and utilities for bidder j when bidding 0, 1, 2, or 3, and when having value $9/4$ or 10:

b	0	1	2	3
$H_j(b)$	$\frac{3}{50n}$	$\frac{3}{50}$	$\frac{53}{100}$	1
$u_j(b; 9/4)$	$\frac{27}{200n}$	$\frac{3}{40}$	$\frac{53}{400}$	–
$u_j(b; 10)$	$\frac{3}{5n}$	$\frac{27}{50}$	$\frac{106}{25}$	7

Since $n \geq 2$, we have the bounds $u_j(0; 9/4) \leq \frac{27}{400}$, $u_j(0; 10) \leq \frac{3}{10}$. Comparing these values, we conclude that it is optimal for bidder j to bid $\hat{\beta}_j(9/4) = 2$, $\hat{\beta}_j(10) = 3$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

4. Suppose bidder i plays the strategy $(0, 3, x)$. We can compute the winning probabilities and utilities for bidder j when bidding 0, 1, 2, or 3, and when having value 9/4 or 10:

b	0	1	2	3
$H_j(b)$	$\frac{3}{50n}$	$\frac{3}{50}$	$\frac{3}{50}$	$\frac{53}{100}$
$u_j(b; 9/4)$	$\frac{27}{200n}$	$\frac{3}{40}$	$\frac{3}{200}$	—
$u_j(b; 10)$	$\frac{3}{5n}$	$\frac{27}{50}$	$\frac{24}{50}$	$\frac{371}{100}$

Since $n \geq 2$, we have the bounds $u_j(0; 9/4) \leq \frac{27}{400}$, $u_j(0; 10) \leq \frac{3}{10}$. Comparing these values, we conclude that it is optimal for bidder j to bid $\hat{\beta}_j(9/4) = 1$, $\hat{\beta}_j(10) = 3$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

From the above case analysis, we obtain the following lemma:

Lemma C.1. *If the bidder i corresponding to the input of a projection gadget plays either s_0 or s_1 , then, in all ε -PBNE of the resulting DFPA (where $\varepsilon \in [0, \frac{1}{18})$), bidder j (corresponding to the gadget's output) matches i 's strategy. Additionally, if bidder i plays some strategy of the form $(0, 1, x)$ or $(0, 2, x)$, bidder j 's unique ε -best-response for $\varepsilon \in [0, \frac{1}{18})$ is s_0 and s_1 respectively.*

Using the construction of the projection gadget, we can similarly define a *split* gadget. Let i be the bidder corresponding to the input of the split, and j, k be the bidders corresponding to its outputs. We can then define the subjective priors of j and k as if each of them was the output of a projection gadget from i . It follows directly from our analysis above that in all ε -PBNE of the resulting DFPA, the output bidders of a split in the circuit have to match the strategy of the input bidder. Additionally, as these splits are used in order to repeat a value in multiple parts of the circuit, it is important to note that the number of bidders introduced remains polynomial on the number of gates of the circuit, since all outputs of a split will end up being inputs to some gate.

C.1.2 OR gadget

Next, we show how to simulate an OR gate. Suppose bidders i, j encode boolean values via their strategies $\hat{\beta}_i$ and $\hat{\beta}_j$. Consider a bidder k whose subjective prior is given as follows: F_{ki} and F_{kj} are discrete distributions that assume value 0 with probability 6/100 and value 9/4 with probability 94/100; for any other $i' \neq i, j$, $F_{ki'}$ is the discrete distribution that assumes value 0 with probability 1. Again, we assume that bidders other than k adopt one of the two strategies $(0, 1, 2)$ or $(0, 2, 3)$. In particular, they bid 0 at value 0, and either 1 or 2 at value 9/4; for our analysis in this part, what they bid at value 10 does not make a difference, as k 's subjective prior puts 0 probability mass at value 10. We can show that bidder k 's best-response when having value 9/4 is to copy the highest among the strategies of bidders i and j . Then, we can add an extra *projection* gadget, mapping k 's strategy which will be of the form either $(0, 1, x)$ or $(0, 2, x')$ back to s_0 and s_1 respectively. Once again, note that k 's best-response must satisfy $\hat{\beta}_k(0) = 0$, due to the no-overbidding assumption. In all the following tables we omit the row corresponding to $u_k(b; 10)$, as this will be mapped back to one of s_0, s_1 with the use of the *projection* gadget.

1. Suppose bidders i and j both play the strategy $(0, 1, 2)$. We can compute the winning probabilities and utilities for bidder k when bidding 0, 1, 2, or 3, and when having value 9/4 or 10:

b	0	1	2	3
$H_k(b)$	$\frac{9}{2500n}$	$\frac{2659}{7500}$	1	1
$u_k(b; 9/4)$	$\frac{81}{10000n}$	$\frac{2659}{6000}$	$\frac{1}{4}$	–

Since $n \geq 3$, we have the bound $u_k(0; 9/4) \leq \frac{27}{10000}$. Comparing these values, we conclude that it is optimal for bidder j to bid $\hat{\beta}_k(9/4) = 1$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

- Suppose bidder i plays the strategy (0, 1, 2) and bidder j plays the strategy (0, 2, 3) (the same analysis holds if we swap the roles of i and j). We can compute the winning probabilities and utilities for bidder k when bidding 0, 1, 2, or 3, and when having value 9/4 or 10:

b	0	1	2	3
$H_k(b)$	$\frac{9}{2500n}$	$\frac{159}{5000}$	$\frac{53}{100}$	1
$u_k(b; 9/4)$	$\frac{81}{10000n}$	$\frac{159}{4000}$	$\frac{53}{400}$	–

Since $n \geq 3$, we have the bound $u_k(0; 9/4) \leq \frac{27}{10000}$. Comparing these values, we conclude that it is optimal for bidder j to bid $\hat{\beta}_k(9/4) = 2$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

- Suppose bidders i and j both play the strategy (0, 2, 3). We can compute the winning probabilities and utilities for bidder k when bidding 0, 1, 2, or 3, and when having value 9/4 or 10:

b	0	1	2	3
$H_k(b)$	$\frac{9}{2500n}$	$\frac{9}{2500}$	$\frac{2659}{7500}$	1
$u_k(b; 9/4)$	$\frac{81}{10000n}$	$\frac{9}{2000}$	$\frac{2659}{30000}$	–

Since $n \geq 3$, we have the bound $u_k(0; 9/4) \leq \frac{27}{10000}$. Comparing these values, we conclude that it is optimal for bidder j to bid $\hat{\beta}_k(9/4) = 2$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

Using the construction of the OR gadget followed by a *projection* gadget, we get the following lemma:

Lemma C.2. *Let i and j be the bidders corresponding to the inputs of an OR gate and assume each of them plays either s_0 or s_1 . Additionally, let k be the bidder introduced in the OR gadget, which we pass as an input to a projection gadget with corresponding output bidder k' . Then, the only ε -best-response of bidder k' (where $\varepsilon \in [0, \frac{1}{18})$) is s_1 , if and only if at least one of i and j plays s_1 , else it is s_0 .*

The above lemma implies that the OR gate is correctly simulated at all PBNE.

C.1.3 Negation gadget

To conclude this section, we show how to simulate a negation gate. Suppose bidder i encodes a boolean value via her strategy $\hat{\beta}_i$. Consider a bidder j whose subjective prior is given as follows: F_{ji} is the discrete distribution that assumes value 0 with probability 6/100 and value 10 with probability 94/100; for any other $i' \neq i, j$, $F_{ji'}$ is the discrete distribution that assumes value 0 with probability 1. Again, assuming that every bidder other than j adopts one of the two strategies (0, 1, 2) or (0, 2, 3), we can show

the best-response of bidder j , when having value $9/4$, is to negate the strategy of bidder i . Afterwards, similarly to what we did for the OR gate, we add a *projection* gadget, mapping the strategy of j (which will be either of the form $(0, 1, x)$ or $(0, 2, x)$) back to s_0 and s_1 respectively. Similarly to the description of the OR gadget, it must be the case that j 's best-response satisfies $\hat{\beta}_j(0) = 0$, due to the no-overbidding assumption. Once again, we omit the row of the table corresponding to $u_j(b; 10)$.

1. Suppose bidder i plays the strategy $(0, 1, 2)$. We can compute the winning probabilities and utilities for bidder j when bidding 0, 1, 2, or 3, and when having a value of $9/4$ or 10:

b	0	1	2	3
$H_j(b)$	$\frac{3}{50n}$	$\frac{3}{50}$	$\frac{53}{100}$	1
$u_j(b; 9/4)$	$\frac{27}{200n}$	$\frac{3}{40}$	$\frac{53}{400}$	—

Since $n \geq 2$, we have the bound $u_j(0; 9/4) \leq \frac{27}{400}$. Comparing these values, we conclude that it is optimal for bidder j to bid $\hat{\beta}_j(9/4) = 2$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

2. Suppose bidder i plays the strategy $(0, 2, 3)$. We can compute the winning probabilities and utilities for bidder j when bidding 0, 1, 2, or 3, and when having a value of $9/4$ or 10:

b	0	1	2	3
$H_j(b)$	$\frac{3}{50n}$	$\frac{3}{50}$	$\frac{3}{50}$	$\frac{53}{100}$
$u_j(b; 9/4)$	$\frac{27}{200n}$	$\frac{3}{40}$	$\frac{3}{200}$	—

Since $n \geq 2$, we have the bound $u_j(0; 9/4) \leq \frac{27}{400}$. Comparing these values, we conclude that it is optimal for bidder j to bid $\hat{\beta}_j(9/4) = 1$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

Let j' be the bidder corresponding to the output of a *projection* gadget, which takes j as input. Again, assuming that every bidder other than j, j' adopts one of the two strategies $(0, 1, 2)$ or $(0, 2, 3)$, we can show the best-response of bidder j' is to negate the strategy of bidder i :

Lemma C.3. *If the bidder i corresponding to the input of a negation gadget plays either s_0 or s_1 , then, in all ε -PBNE of the resulting DFPA (where $\varepsilon \in [0, \frac{1}{18})$), bidder j' (corresponding to the output of the projection gadget added after the negation gate) plays the opposite of i 's strategy.*

C.2 Input gadget

We can see (and trivially prove by structural induction on the circuit) that, given that the bidders corresponding to the inputs of the circuit play either s_0 or s_1 , all bidders in the DFPA are playing ε -best-responses if and only if the gates are correctly simulated. However, it remains to specify the subjective priors for these input bidders. To achieve this normalized behaviour, we can introduce a special gadget, similar to the *projection* gadget, for each input to the circuit. Essentially, we add two bidders projecting to each other; our analysis will show that in all ε -PBNE they should both be playing either s_0 or s_1 .

Suppose bidder i encodes a boolean value via her strategy $\hat{\beta}_i$. Consider a bidder j , whose subjective prior is given as follows: F_{ji} is the discrete distribution that assumes value 0 with probability $6/100$ and value $9/4$ with probability $94/100$; for any other $i' \neq i, j$, $F_{ji'}$ is the discrete distribution that

assumes value 0 with probability 1. At the same time, i 's subjective prior is given by: F_{ij} is the discrete distribution that assumes value 0 with probability $6/100$ and value $9/4$ with probability $94/100$; for any other $j' \neq i, j$, $F_{ij'}$ is the discrete distribution that assumes value 0 with probability 1.

Using the analysis from [Appendix C.1.1](#), we can list all the best-responses of j in a potential PBNE:

i 's strategy	j 's best-response
(0, 0, x)	(0, 1, 1)
(0, 1, x)	(0, 1, 2)
(0, 2, x)	(0, 2, 3)
(0, 3, x)	(0, 1, 3)

Table 1: Bidder j 's best-responses given i 's strategies

Also, from previous sections, we know that there are no other strategies that result in utility within $1/18$ of that of the best-responses. It is evident that, due to the symmetric nature of the input gadget, the best-responses of i to the strategies of j are symmetrical to the ones in the above table. Therefore, in all PBNE i and j would have to play the same strategy, which would be either s_0 or s_1 , thus yielding the intended behaviour of the circuit's input gates.

Lemma C.4. *If i and j are the bidders corresponding to some input of the circuit (as described in the construction of the input gadget), then, in all ε -PBNE for $\varepsilon \in [0, 1/18)$, i and j both play the same strategy, which can be either s_0 or s_1 .*

Combining [Lemmas C.1](#) to [C.4](#), we deduce the following:

Corollary C.5. *Using the above mapping from a circuit to a DFPA, all induced bidders are simultaneously ε -best-responding (for $\varepsilon \in [0, 1/18)$) to each other if and only if the values that their strategies map to satisfy the gates of the circuit.*

C.3 Output Gadget

Thus far, our construction ensures that at an ε -PBNE it is necessary that all the gates are correctly simulated. As we are reducing from the CIRCUIT-SAT problem, we want an ε -PBNE to exist only in cases where the output of the circuit can take *true* as a value. This leads us to the last gadget of our construction - the *output* gadget. This introduces two new bidders to the auction, k and ℓ , which (given their subjective priors, which assume value 0 with probability 1 for every bidder other than each other) cannot play best-responses simultaneously. Then, we also add k 's prior distribution for bidder i (the bidder corresponding to the output of the circuit), such that now the bidders can simultaneously play ε -best-responses if and only if bidder i 's strategy represents the value *true*.

Next, we provide the details of the construction of the output gadget. We introduce the two new bidders k and ℓ with the following prior distributions: $F_{k\ell}$ and $F_{\ell k}$ are discrete distributions that assume value 0 with probability $8/11$ and value 10 with probability $3/11$. Bidder k 's prior for bidder i (corresponding to the circuit's output) is F_{ki} , the discrete distribution that assumes value 0 with probability $1/2$ and value $9/4$ with probability $1/2$. For any other $k' \neq k, \ell$ and $\ell' \neq k, \ell, i$, $F_{\ell k'}$ and $F_{k\ell'}$ are discrete distributions that assume value 0 with probability 1.

We will begin our analysis by calculating bidder ℓ 's best-responses depending on k 's strategies. Note that, since $F_{\ell k}$ only has mass at 0 and 10, it doesn't matter what $\hat{\beta}_k(9/4)$ is for strategy $\hat{\beta}_k$, hence we will denote this as x in our analysis. As our aim is to compare how k and ℓ best-respond to each other, and the subjective prior of each puts positive probability only on 0 and 10, we will only analyse what $u_\ell(b; 10)$ will be. We begin by noting that all of ℓ 's best-responses must satisfy $\hat{\beta}_\ell(0) = 0$, due to the no-overbidding assumption. Then:

1. Suppose bidder k plays some strategy of the form $(0, x, 0)$. We can compute the winning probabilities and utilities for bidder ℓ when bidding 0, 1, 2, or 3 and having value 10:

b	0	1	2	3
$H_\ell(b)$	$\frac{1}{n}$	1	1	1
$u_\ell(b; 10)$	$\frac{10}{n}$	9	8	7

Since $n \geq 2$, we have the bound $u_\ell(0; 10) \leq 5$. Comparing the values in the table, we conclude that it is optimal for bidder ℓ to bid $\hat{\beta}_\ell(10) = 1$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

2. Suppose bidder k plays some strategy of the form $(0, x, 1)$. We can compute the winning probabilities and utilities for bidder ℓ when bidding 0, 1, 2, or 3 and having value 10:

b	0	1	2	3
$H_\ell(b)$	$\frac{8}{11n}$	$\frac{19}{22}$	1	1
$u_\ell(b; 10)$	$\frac{80}{11n}$	$\frac{171}{22}$	8	7

Since $n \geq 2$, we have the bound $u_\ell(0; 10) \leq \frac{40}{11}$. Comparing the values in the table, we conclude that it is optimal for bidder ℓ to bid $\hat{\beta}_\ell(10) = 2$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

3. Suppose bidder k plays some strategy of the form $(0, x, 2)$. We can compute the winning probabilities and utilities for bidder ℓ when bidding 0, 1, 2, or 3 and having value 10:

b	0	1	2	3
$H_\ell(b)$	$\frac{8}{11n}$	$\frac{8}{11}$	$\frac{19}{22}$	1
$u_\ell(b; 10)$	$\frac{80}{11n}$	$\frac{72}{11}$	$\frac{76}{11}$	7

Since $n \geq 2$, we have the bound $u_\ell(0; 10) \leq \frac{40}{11}$. Comparing these values, we conclude that it is optimal for bidder ℓ to bid $\hat{\beta}_\ell(9/4) = 1$, $\hat{\beta}_\ell(10) = 3$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

4. Suppose bidder k plays some strategy of the form $(0, x, 3)$. We can compute the winning probabilities and utilities for bidder ℓ when bidding 0, 1, 2, or 3 and having value 10:

b	0	1	2	3
$H_\ell(b)$	$\frac{8}{11n}$	$\frac{8}{11}$	$\frac{8}{11}$	$\frac{19}{22}$
$u_\ell(b; 10)$	$\frac{80}{11n}$	$\frac{72}{11}$	$\frac{64}{11}$	$\frac{133}{22}$

Since $n \geq 2$, we have the bound $u_\ell(0; 10) \leq \frac{40}{11}$. Comparing these values, we conclude that it is optimal for bidder ℓ to bid $\hat{\beta}_\ell(10) = 1$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

We can note that k 's strategy at an equilibrium (from ℓ 's point of view) can be represented by just $\hat{\beta}_k(10)$, since $\hat{\beta}_k(0) = 0$ at all PBNE due to the no-overbidding assumption. We can do the same thing for ℓ , since k 's prior for her has mass at a point other than 0, and that is at 10, so we get the following table representing ℓ 's best-responses to k :

$\hat{\beta}_k(10)$	0	1	2	3
$\hat{\beta}_\ell(10)$	1	2	3	1

Table 2: Bidder ℓ 's best-responses to bidder k 's strategies

We will now look at k 's best-response. Since k has a non-zero prior for both i and ℓ , her best-response will depend on two other bidders, resulting in more cases. We begin our analysis with the ones where i plays $(0, 1, 2)$. Our goal is to show that none of them can lead to a PBNE of the DFPA. It is also safe to assume that, at all PBNE, i plays either $(0, 1, 2)$ or $(0, 2, 3)$, using the analysis of the simulation of the circuit's gates, as i corresponds to the output of the circuit. Start by fixing i 's strategy to $s_0 = (0, 1, 2)$. By the no-overbidding assumption, k 's best-response must satisfy $\hat{\beta}_k(0) = 0$. Then, we consider the following cases:

1. Suppose bidder ℓ plays some strategy of the form $(0, x, 0)$. We can compute the winning probabilities and utilities for bidder k when bidding 0, 1, 2, or 3 and having value 10:

b	0	1	2	3
$H_k(b)$	$\frac{1}{2n}$	$\frac{3}{4}$	1	1
$u_k(b; 10)$	$\frac{5}{n}$	$\frac{27}{4}$	8	7

Since $n \geq 3$, we have the bound $u_k(0; 10) \leq \frac{5}{3}$. Comparing the values in the table, we conclude that it is optimal for bidder k to bid $\hat{\beta}_k(10) = 2$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

2. Suppose bidder ℓ plays some strategy of the form $(0, x, 1)$. We can compute the winning probabilities and utilities for bidder k when bidding 0, 1, 2, or 3 and having value 10:

b	0	1	2	3
$H_k(b)$	$\frac{4}{11n}$	$\frac{29}{44}$	1	1
$u_k(b; 10)$	$\frac{40}{11n}$	$\frac{261}{44}$	8	7

Since $n \geq 3$, we have the bound $u_k(0; 10) \leq \frac{40}{33}$. Comparing the values in the table, we conclude that it is optimal for bidder k to bid $\hat{\beta}_k(10) = 2$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

3. Suppose bidder ℓ plays some strategy of the form $(0, x, 2)$. We can compute the winning probabilities and utilities for bidder k when bidding 0, 1, 2, or 3 and having value 10:

b	0	1	2	3
$H_k(b)$	$\frac{4}{11n}$	$\frac{6}{11}$	$\frac{19}{22}$	1
$u_k(b; 10)$	$\frac{40}{11n}$	$\frac{72}{11}$	$\frac{76}{11}$	7

Since $n \geq 3$, we have the bound $u_k(0; 10) \leq \frac{40}{33}$. Comparing the values in the table, we conclude that it is optimal for bidder k to bid $\hat{\beta}_k(10) = 3$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

4. Suppose bidder ℓ plays some strategy of the form $(0, x, 3)$. We can compute the winning probabilities and utilities for bidder k when bidding 0, 1, 2, or 3 and having value 10:

b	0	1	2	3
$H_k(b)$	$\frac{4}{11n}$	$\frac{6}{11}$	$\frac{8}{11}$	$\frac{19}{22}$
$u_k(b; 10)$	$\frac{40}{11n}$	$\frac{54}{11}$	$\frac{64}{11}$	$\frac{133}{22}$

Since $n \geq 3$, we have the bound $u_k(0; 10) \leq \frac{65}{33}$. Comparing the values in the table, we conclude that it is optimal for bidder k to bid $\hat{\beta}_k(10) = 3$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$.

Hence, k 's unique (ε -) best-responses for $\varepsilon < \frac{1}{18}$ to ℓ 's strategies when i plays $(0, 1, 2)$ are the ones demonstrated in the following table:

$\hat{\beta}_\ell(10)$	0	1	2	3
$\hat{\beta}_k(10)$	2	2	3	3

Table 3: Bidder k 's best-responses given ℓ 's strategies when i plays $(0, 1, 2)$

Comparing this table with the one describing ℓ 's best-responses, we can easily see that there is no equilibrium when i plays $(0, 1, 2)$. We now examine the case where i plays $(0, 2, 3)$. It suffices to show that there is some strategy of ℓ for which k 's best-response leads to both playing best-responses simultaneously. To this end, we consider the case where ℓ plays some strategy $(0, x, 1)$. Then, k 's utility for each possible bid is the following:

b	0	1	2	3
$H_k(b)$	$\frac{4}{11n}$	$\frac{19}{44}$	$\frac{3}{4}$	1
$u_k(b; 10)$	$\frac{40}{11n}$	$\frac{171}{44}$	6	7

Since $n \geq 3$, we have the bound $u_k(0; 10) \leq \frac{40}{33}$. Comparing the values in the table, we conclude that it is optimal for bidder k to bid $\hat{\beta}_k(10) = 3$. Also, we can see that there is no other ε -best-response for $\varepsilon < \frac{1}{18}$. Consequently, there is a strategy profile with k and ℓ simultaneously playing best-responses, where ℓ plays some strategy $(0, x, 1)$ and k plays some strategy $(0, x', 3)$. This concludes our proof of the following lemma:

Lemma C.6. *Given the construction of the output gadget, agents k and ℓ can simultaneously ε -best-respond to every agent (for $\varepsilon \in [0, 1/18)$) if and only if the bidder corresponding to the output of the circuit plays the strategy $s_1 = (0, 2, 3)$.*

Combining the results of [Appendices C.1 to C.3](#), we complete the proof of the NP-hardness result. Note that this will hold for $1/10$ of the value of ε we used in the reduction, due to the scaling argument that we mentioned in the beginning of [Appendix C](#). We will demonstrate that there is a satisfying assignment to the circuit's inputs if and only if there is a PBNE in the resulting DFPA.

Assume the circuit is satisfiable and let α be a satisfying assignment to its inputs. Then, using [Corollary C.5](#) and [Lemma C.6](#), the resulting DFPA has an exact PBNE where for each input α_i to the circuit, the corresponding bidders i and j (as described in [Appendix C.2](#)) both play strategy s_0 if $\alpha_i = \text{false}$, and s_1 otherwise.

If on the other hand the circuit is unsatisfiable, then the output bidder would always play s_0 in order to best-respond to the previous bidders while satisfying the circuit's gates, so there can be no ε -PBNE in the resulting DFPA, due to [Lemma C.6](#). This concludes our proof of NP-hardness in the auction we introduced for $\varepsilon < \frac{1}{18}$, thus yielding NP-hardness for $\varepsilon < \frac{1}{180}$ in the original auction.

D PPAD-Hardness of Finding Mixed Equilibria

In this section, we prove the following result:

Theorem D.1. *For any $\varepsilon < 1/36$, it is PPAD-hard to compute an ε -well-supported MBNE in a DFPA, even when the auction has interaction degree bounded by $d = 2$.*

By Lemma 4.1 we also obtain the following.

Corollary D.2. *For any $\varepsilon < 1/(3^4 \cdot 2^8)$, it is PPAD-hard to compute an ε -approximate MBNE in a DFPA.*

D.1 The PURE-CIRCUIT Problem

An instance of the PURE-CIRCUIT problem is given by a node set X and a set G of gates. Each gate $g \in G$ is of the form $g = (T, x, y, z)$ where $x, y, z \in X$ are distinct nodes, and $T \in \{\text{NOT}, \text{AND}, \text{PURIFY}\}$ is the type of the gate, with the following interpretation.

- If $T = \text{NOT}$, then x is the input of the gate, and y is its output. (z is unused)
- If $T = \text{AND}$, then x and y are the inputs of the gate, and z is its output.
- If $T = \text{PURIFY}$, then x is the input of the gate, and y and z are its outputs.

We require that each node is the output of exactly one gate.

A solution to instance (X, G) is an assignment $\mathbf{A} : X \rightarrow \{0, 1, \perp\}$ that satisfies all the gates, i.e., for each gate $g = (T, x, y, z) \in G$ we have the following.

- If $T = \text{NOT}$ in $g = (T, x, y)$, then \mathbf{A} satisfies

$$\begin{aligned} \mathbf{A}[x] = 0 &\implies \mathbf{A}[y] = 1 \\ \mathbf{A}[x] = 1 &\implies \mathbf{A}[y] = 0. \end{aligned}$$

- If $T = \text{AND}$ in $g = (T, x, y, z)$, then \mathbf{A} satisfies

$$\begin{aligned} \mathbf{A}[x] = \mathbf{A}[y] = 1 &\implies \mathbf{A}[z] = 1 \\ \mathbf{A}[x] = 0 \vee \mathbf{A}[y] = 0 &\implies \mathbf{A}[z] = 0. \end{aligned}$$

- If $T = \text{PURIFY}$, then \mathbf{A} satisfies

$$\begin{aligned} \{\mathbf{A}[y], \mathbf{A}[z]\} \cap \{0, 1\} &\neq \emptyset \\ \mathbf{A}[x] \in \{0, 1\} &\implies \mathbf{A}[y] = \mathbf{A}[z] = \mathbf{A}[x]. \end{aligned}$$

Theorem D.3 ([DFHM22]). *The PURE-CIRCUIT problem is PPAD-complete.*

D.2 Construction of the Auction

Consider an instance (X, G) of PURE-CIRCUIT. We now describe how to construct an instance of the discrete FPA that encodes the PURE-CIRCUIT instance. In the next section, we will then argue that any approximate equilibrium of the auction must yield a solution to (X, G) .

The auction consists of one bidder for each of the variables in X , along with a number of additional auxiliary bidders. Each bidder i has a value space consisting of two values, $V_i = \{0, v_i\}$, where $v_i \in (0, 1]$

depends on i . The bidding space is $B = \{b_0, b_1, b_2, b_3\} = \{0, 1/4, 1/2, 3/4\}$. We now describe the bidders in the auction.

Constant bidders. The purpose of this part of the construction is to have a bidder const who bids b_2 with probability 1 when its value is v_{const} in any approximate equilibrium. In order to achieve this, we introduce a set C of two additional bidders and ensure that they always bid b_1 when their value is non-zero. In more detail, each bidder $i \in C$ has value space $V_i = \{0, v_i\}$ with $v_i = 1/2$, and we set $f_{i,j}(0) = 1, f_{i,j}(v_j) = 0$ for all other bidders j in the auction. The bidder const has value space $V_{\text{const}} = \{0, v_{\text{const}}\}$ with $v_{\text{const}} = 1$, and we set $f_{\text{const},j}(0) = 0, f_{\text{const},j}(v_j) = 1$ for all $j \in C$, and $f_{i,j}(0) = 1, f_{i,j}(v_j) = 0$ for all other bidders j in the auction. We show that this construction indeed satisfies the desired properties in the next section.

Variable bidders. For every variable $x \in X$ of the PURE-CIRCUIT instance, we introduce a corresponding bidder x in the auction. The value space V_x and distributions $f_{x,j}$ depend on the type of gate that uses x as output.

AND gate. For every variable $z \in X$ that is the output of an AND gate with inputs x and y , we set the value space to be $V_z = \{0, v_z\}$, where $v_z = 7/12$. The distributions are $f_{z,j}(0) = 0, f_{z,j}(v_j) = 1$ for $j \in \{x, y\}$, and $f_{z,j}(0) = 1, f_{z,j}(v_j) = 0$ for all other bidders j in the auction.

PURIFY gate. For every variable $y \in X$ that is the *first* output of a PURIFY gate with input x , we set the value space to be $V_y = \{0, v_y\}$, where $v_y = 9/16$. The distributions are $f_{y,x}(0) = 0, f_{y,x}(v_x) = 1$, and $f_{y,j}(0) = 1, f_{y,j}(v_j) = 0$ for all other bidders j in the auction.

For every variable $z \in X$ that is the *second* output of a PURIFY gate with input x , we set the value space to be $V_z = \{0, v_z\}$, where $v_z = 11/16$. The distributions are $f_{z,x}(0) = 0, f_{z,x}(v_x) = 1$, and $f_{z,j}(0) = 1, f_{z,j}(v_j) = 0$ for all other bidders j in the auction.

NOT gate. For every variable $y \in X$ that is the output of a NOT gate with input x , we introduce an additional auxiliary bidder y' to help with the implementation of the gate. The value space of bidder y' is $V_{y'} = \{0, v_{y'}\}$, where $v_{y'} = 13/14$. The distributions are $f_{y',x}(0) = 0, f_{y',x}(v_x) = 1, f_{y',\text{const}}(0) = 0, f_{y',\text{const}}(v_{\text{const}}) = 1$, and $f_{y',j}(0) = 1, f_{y',j}(v_j) = 0$ for all other bidders j in the auction.

The value space of bidder y is $V_y = \{0, v_y\}$, where $v_y = 5/8$. The distributions are $f_{y,y'}(0) = 1/9, f_{y,y'}(v_{y'}) = 8/9$, and $f_{y,j}(0) = 1, f_{y,j}(v_j) = 0$ for all other bidders j in the auction.

Properties. This instance of the auction problem can clearly be constructed in polynomial time given an instance (X, G) of PURE-CIRCUIT. Furthermore, by construction the auction has interaction degree bounded by $d = 2$.

D.3 Analysis

Fix any constant $\varepsilon < 1/36$. Without loss of generality, we can assume that $|X|$ is sufficiently large so that the number of bidders in the auction, denoted by n , satisfies $n \geq \max\{1000, 1/(1/36 - \varepsilon)\}$.

Consider any ε -well-supported equilibrium β of the auction. In this section, we show how to extract a solution to the PURE-CIRCUIT instance (X, G) from β in polynomial time.

Extraction. We construct an assignment $\mathbf{A} : X \rightarrow \{0, 1, \perp\}$ to the variables of PURE-CIRCUIT by letting

- $\mathbf{A}[x] = 1$ if $\beta_x(v_x)(b_1) = 1$,
- $\mathbf{A}[x] = 0$ if $\beta_x(v_x)(b_1) = 0$,
- $\mathbf{A}[x] = \perp$ otherwise.

In the remainder of this section, we argue that this assignment satisfies all the gates of the PURE-CIRCUIT instance (X, G) . Note that by construction we have $v_x < b_3$ for all $x \in X$, so we will necessarily have

$\beta_x(v_x)(b_3) = 0$. For a bidder $x \in X$, we say that x is *valid*, if, in addition, $\beta_x(v_x)(b_0) = 0$, or, in other words,

$$\beta_x(v_x)(b_1) + \beta_x(v_x)(b_2) = 1.$$

Ultimately, we will argue that all bidders $x \in X$ must be valid.

Constant bidders. We start by showing that the constant bidders have the desired behaviour.

Lemma D.4. *We have $\beta_{\text{const}}(v_{\text{const}})(b_2) = 1$.*

Proof. We first analyse the behaviour of the bidders in set C . Consider any bidder $i \in C$ and note that since $f_{i,j}(0) = 1$ for all $j \neq i$, bidder i believes that all other bidders have value 0. As a result, by the no-overbidding assumption, bidder i believes that all other bidders will bid $b_0 = 0$. This yields the following utilities for bidder i at value $v_i = 1/2$:

$$\begin{aligned} u_i(b_0, \beta_{-i}; v_i) &= (v_i - b_0) \cdot H_i(b_0, \beta_{-i}) = \frac{1}{2} \cdot \frac{1}{n} \\ u_i(b_1, \beta_{-i}; v_i) &= (v_i - b_1) \cdot H_i(b_1, \beta_{-i}) = \frac{1}{4} \\ u_i(b_2, \beta_{-i}; v_i) &= (v_i - b_2) \cdot H_i(b_2, \beta_{-i}) = 0 \end{aligned}$$

and bid b_3 is not playable because $b_3 > v_i$. Since $n \geq 1000$, we have $1/4 > 1/2n + \varepsilon$, and thus b_1 is the only ε -best-response, i.e., $\beta_i(v_i)(b_1) = 1$.

Now we turn our attention to the bidder const . Recall that $f_{\text{const},j}(0) = 0, f_{\text{const},j}(v_j) = 1$ for all $j \in C$, and $f_{i,j}(0) = 1, f_{i,j}(v_j) = 0$ for all other bidders j in the auction. As a result, since $\beta_j(v_j)(b_1) = 1$ for all $j \in C$, bidder const believes that all bidders $j \in C$ bid b_1 . Furthermore, by the same argument as above, bidder const believes that all other bidders outside C bid $b_0 = 0$. This yields the following utilities for bidder $i = \text{const}$ at value $v_i = v_{\text{const}} = 1$:

$$\begin{aligned} u_i(b_0, \beta_{-i}; v_i) &= (v_i - b_0) \cdot H_i(b_0, \beta_{-i}) = 0 \\ u_i(b_1, \beta_{-i}; v_i) &= (v_i - b_1) \cdot H_i(b_1, \beta_{-i}) = \frac{3}{4} \cdot \frac{1}{|C| + 1} \\ u_i(b_2, \beta_{-i}; v_i) &= (v_i - b_2) \cdot H_i(b_2, \beta_{-i}) = \frac{1}{2} \\ u_i(b_3, \beta_{-i}; v_i) &= (v_i - b_3) \cdot H_i(b_3, \beta_{-i}) = \frac{1}{4} \end{aligned}$$

Recalling that $|C| = 2$, we see that b_2 is the only ε -best-response, since $\varepsilon < 1/4$. □

AND gate. Next, we consider bidders corresponding to an AND gate.

Lemma D.5. *For any AND gate with inputs x and y , and output z , we have that*

- z is valid,
- if x and y are valid, then A satisfies the conditions of the AND gate.

Proof. Let us first show that z is valid. By construction, bidder z believes that all other bidders, except x and y , will bid b_0 . Furthermore, z believes that x and y will bid according to $\beta_x(v_x)$ and $\beta_y(v_y)$, respectively. Recalling that $\beta_x(v_x)(b_3) = \beta_y(v_y)(b_3) = 0$, and using the notation $p_j = \beta_x(v_x)(b_j)$, $q_j = \beta_y(v_y)(b_j)$, for $j \in \{0, 1, 2\}$, we can write the utilities of bidder $i = z$ at value $v_i = v_z = 7/12$:

$$\begin{aligned} u_i(b_0, \beta_{-i}; v_i) &= (v_i - b_0) \cdot H_i(b_0, \beta_{-i}) = \frac{7}{12} p_0 q_0 / n \\ u_i(b_1, \beta_{-i}; v_i) &= (v_i - b_1) \cdot H_i(b_1, \beta_{-i}) = \left(\frac{7}{12} - \frac{1}{4} \right) (p_0 q_0 + p_1 q_0 / 2 + p_0 q_1 / 2 + p_1 q_1 / 3) \\ u_i(b_2, \beta_{-i}; v_i) &= \left(\frac{7}{12} - \frac{1}{2} \right) ((1 - p_2)(1 - q_2) + (1 - p_2)q_2 / 2 + p_2(1 - q_2) / 2 + p_2 q_2 / 3) \end{aligned}$$

and bid b_3 is not playable because $b_3 > v_i$. Note that we always have $u_i(b_2, \boldsymbol{\beta}_{-i}; v_i) \geq (7/12 - 1/2)/3 = 1/36$, because $(1-p_2)(1-q_2) + (1-p_2)q_2/2 + p_2(1-q_2)/2 + p_2q_2/3 \geq 1/3$. Furthermore, $u_i(b_0, \boldsymbol{\beta}_{-i}; v_i) \leq 7/12n < 1/n \leq 1/36 - \varepsilon$, since $n \geq 1/(1/36 - \varepsilon)$. As a result, b_0 can never be an ε -best-response for z , and thus z is valid.

Now consider the case where x and y are valid, i.e., $p_0 = q_0 = 0$. In that case, the utilities can be simplified to

$$\begin{aligned} u_i(b_0, \boldsymbol{\beta}_{-i}; v_i) &= 0 \\ u_i(b_1, \boldsymbol{\beta}_{-i}; v_i) &= \left(\frac{7}{12} - \frac{1}{4}\right) p_1 q_1 / 3 \\ u_i(b_2, \boldsymbol{\beta}_{-i}; v_i) &= \left(\frac{7}{12} - \frac{1}{2}\right) (p_1 q_1 + p_1(1-q_1)/2 + (1-p_1)q_1/2 + (1-p_1)(1-q_1)/3) \end{aligned}$$

If $\mathbf{A}[x] = \mathbf{A}[y] = 1$, i.e., $p_1 = q_1 = 1$, then

$$u_i(b_1, \boldsymbol{\beta}_{-i}; v_i) - u_i(b_2, \boldsymbol{\beta}_{-i}; v_i) = \left(\frac{7}{12} - \frac{1}{4}\right) \frac{1}{3} - \left(\frac{7}{12} - \frac{1}{2}\right) = \frac{1}{36} > \varepsilon$$

and so b_1 is the only ε -best-response, i.e., $\mathbf{A}[z] = 1$.

Finally, consider the case where $\mathbf{A}[x] = 0$ or $\mathbf{A}[y] = 0$. Without loss of generality, given the symmetry of the construction, assume that $\mathbf{A}[x] = 0$, i.e., $p_1 = 0$. In that case, we have

$$u_i(b_2, \boldsymbol{\beta}_{-i}; v_i) - u_i(b_1, \boldsymbol{\beta}_{-i}; v_i) = \frac{1}{12}(q_1/2 + (1-q_1)/3) - 0 \geq \frac{1}{36} > \varepsilon$$

and so b_2 is the only ε -best-response, i.e., $\mathbf{A}[z] = 0$. □

PURIFY gate. Next, we consider bidders corresponding to a PURIFY gate.

Lemma D.6. *For any PURIFY gate with input x and outputs y and z , we have that*

- y and z are valid,
- if x is valid, then \mathbf{A} satisfies the conditions of the PURIFY gate.

Proof. Let us first show that y and z are valid. By construction, bidder y believes that all other bidders except x will bid b_0 . Furthermore, y believes that x will bid according to $\beta_x(v_x)$. Recalling that $\beta_x(v_x)(b_3) = 0$, and using the notation $p_j = \beta_x(v_x)(b_j)$ for $j \in \{0, 1, 2\}$, we can write the utilities of bidder $i = y$ at value $v_i = v_y = 9/16$:

$$\begin{aligned} u_i(b_0, \boldsymbol{\beta}_{-i}; v_i) &= (v_i - b_0) \cdot H_i(b_0, \boldsymbol{\beta}_{-i}) = \frac{9}{16} p_0 / n \\ u_i(b_1, \boldsymbol{\beta}_{-i}; v_i) &= (v_i - b_1) \cdot H_i(b_1, \boldsymbol{\beta}_{-i}) = \left(\frac{9}{16} - \frac{1}{4}\right) (p_0 + p_1/2) \\ u_i(b_2, \boldsymbol{\beta}_{-i}; v_i) &= (v_i - b_2) \cdot H_i(b_2, \boldsymbol{\beta}_{-i}) = \left(\frac{9}{16} - \frac{1}{2}\right) (p_0 + p_1 + p_2/2) \end{aligned}$$

and bid b_3 is not playable because $b_3 > v_i$. Since $p_0 + p_1 + p_2/2 \geq 1/2$, we can write

$$u_i(b_2, \boldsymbol{\beta}_{-i}; v_i) - u_i(b_0, \boldsymbol{\beta}_{-i}; v_i) \geq \left(\frac{9}{16} - \frac{1}{2}\right) \frac{1}{2} - \frac{9}{16} p_0 / n = \frac{1}{32} - \frac{9}{16} p_0 / n > \varepsilon$$

because $n \geq 1000$. As a result, b_0 can never be an ε -best-response for y , and thus y is valid. The same reasoning for $i = z$, with $v_i = v_z = 11/16$, yields

$$u_i(b_2, \beta_{-i}; v_i) - u_i(b_0, \beta_{-i}; v_i) \geq \left(\frac{11}{16} - \frac{1}{2}\right) \frac{1}{2} - \frac{11}{16} p_0/n = \frac{3}{32} - \frac{11}{16} p_0/n > \varepsilon$$

because $n \geq 1000$. Thus, b_0 can never be an ε -best-response for z , and z is also valid.

Now consider the case where x is valid, i.e., $p_0 = 0$. In that case, the utilities for $i \in \{y, z\}$ at value v_i can be simplified to

$$\begin{aligned} u_i(b_0, \beta_{-i}; v_i) &= 0 \\ u_i(b_1, \beta_{-i}; v_i) &= \left(v_i - \frac{1}{4}\right) p_1/2 \\ u_i(b_2, \beta_{-i}; v_i) &= \left(v_i - \frac{1}{2}\right) (1/2 + p_1/2) \end{aligned}$$

We can write

$$\Delta := u_i(b_1, \beta_{-i}; v_i) - u_i(b_2, \beta_{-i}; v_i) = \left(v_i - \frac{1}{4}\right) p_1/2 - \left(v_i - \frac{1}{2}\right) (1/2 + p_1/2) = p_1/8 - v_i/2 + 1/4.$$

Now for $i = y$, and thus $v_i = 9/16$, we obtain that

- if $\mathbf{A}[x] = 0$, i.e., $p_1 = 0$, then $\Delta = p_1/8 - 1/32 = -1/32 < -\varepsilon$, and b_2 is the only ε -best-response, i.e., $\mathbf{A}[y] = 0$.
- if $\mathbf{A}[x] \geq 1/2$, i.e., $p_1 \geq 1/2$, then $\Delta = p_1/8 - 1/32 \geq 1/32 > \varepsilon$, and b_1 is the only ε -best-response, i.e., $\mathbf{A}[y] = 1$.

Similarly, for $i = z$, and thus $v_i = 11/16$, we obtain that

- if $\mathbf{A}[x] \leq 1/2$, i.e., $p_1 \leq 1/2$, then $\Delta = p_1/8 - 3/32 \leq -1/32 < -\varepsilon$, and b_2 is the only ε -best-response, i.e., $\mathbf{A}[z] = 0$.
- if $\mathbf{A}[x] = 1$, i.e., $p_1 = 1$, then $\Delta = p_1/8 - 3/32 = 1/32 > \varepsilon$, and b_1 is the only ε -best-response, i.e., $\mathbf{A}[z] = 1$.

As a result, the conditions of the PURIFY gate are indeed satisfied. In particular, we always have that at least one of $\mathbf{A}[y]$ or $\mathbf{A}[z]$ lies in $\{0, 1\}$. \square

NOT gate. Finally, we consider bidders corresponding to a NOT gate.

Lemma D.7. *For any NOT gate with input x and output y , we have that*

- y is valid,
- if x is valid, then \mathbf{A} satisfies the conditions of the NOT gate.

Proof. We first show that y is valid. By construction, bidder y believes that all other bidders except y' will bid b_0 . Furthermore, y believes that y' will bid according to $\beta_{y'}(0)$ with probability $1/9$ (namely when y' has value 0), and according to $\beta_{y'}(v_{y'})$ with probability $8/9$ (namely when y' has value $v_{y'}$).

By the no-overbidding assumption, $\beta_{y'}(0)$ is the distribution with $\beta_{y'}(0)(b_0) = 1$. Using the notation $q_j = \beta_{y'}(v_{y'})(b_j)$ for $j = 0, 1, 2, 3$, we can write the utilities of bidder $i = y$ at value $v_i = v_y = 5/8$:

$$\begin{aligned} u_i(b_0, \boldsymbol{\beta}_{-i}; v_i) &= \frac{5}{8} \left(\frac{1}{9} \cdot \frac{1}{n} + \frac{8}{9} \cdot \frac{q_0}{n} \right) \\ u_i(b_1, \boldsymbol{\beta}_{-i}; v_i) &= \left(\frac{5}{8} - \frac{1}{4} \right) \left(\frac{1}{9} + \frac{8}{9} (q_0 + q_1/2) \right) \\ u_i(b_2, \boldsymbol{\beta}_{-i}; v_i) &= \left(\frac{5}{8} - \frac{1}{2} \right) \left(\frac{1}{9} + \frac{8}{9} (q_0 + q_1 + q_2/2) \right) \end{aligned}$$

and bid b_3 is not playable because $b_3 > v_i$. We always have

$$u_i(b_1, \boldsymbol{\beta}_{-i}; v_i) - u_i(b_0, \boldsymbol{\beta}_{-i}; v_i) \geq \frac{3}{8} \cdot \frac{1}{9} - \frac{5}{8} \cdot \left(\frac{1}{9} \cdot \frac{1}{n} + \frac{8}{9} \cdot \frac{1}{n} \right) = \frac{3}{72} - \frac{5}{8n} > \varepsilon$$

because $n \geq 1000$. As a result, b_0 is never an ε -best-response for y at value v_y , and thus y is valid. Furthermore, we note the following, which will be useful below

– if $\beta_{y'}(v_{y'})(b_2) = 1$, i.e., $q_2 = 1$, then

$$u_i(b_2, \boldsymbol{\beta}_{-i}; v_i) - u_i(b_1, \boldsymbol{\beta}_{-i}; v_i) = \frac{1}{8} \left(\frac{1}{9} + \frac{8}{9} \cdot \frac{1}{2} \right) - \frac{3}{8} \cdot \frac{1}{9} = \frac{1}{36} > \varepsilon$$

and thus b_2 is the only ε -best-response, i.e., $A[y] = 0$.

– if $\beta_{y'}(v_{y'})(b_3) = 1$, i.e., $q_3 = 1$, then

$$u_i(b_1, \boldsymbol{\beta}_{-i}; v_i) - u_i(b_2, \boldsymbol{\beta}_{-i}; v_i) = \frac{3}{8} \cdot \frac{1}{9} - \frac{1}{8} \cdot \frac{1}{9} = \frac{1}{36} > \varepsilon$$

and thus b_1 is the only ε -best-response, i.e., $A[y] = 1$.

From now on, we assume that x is valid, and aim to show that the constraints of the NOT gate are satisfied. By construction, bidder y' believes that all other bidders except x and const will bid b_0 . By [Lemma D.4](#) we know that bidder y' believes that bidder const will bid b_2 . Finally, bidder y' believes that x will bid according to $\beta_x(v_x)$, where $\beta_x(v_x)(b_0) = \beta_x(v_x)(b_3) = 0$, since x is valid. Using the notation $p_1 = \beta_x(v_x)(b_1)$, and thus $1 - p_1 = \beta_x(v_x)(b_2)$, we can write the utilities of bidder $i = y'$ at value $v_i = v_{y'} = 13/14$:

$$\begin{aligned} u_i(b_0, \boldsymbol{\beta}_{-i}; v_i) &= \frac{13}{14} \cdot 0 = 0 \\ u_i(b_1, \boldsymbol{\beta}_{-i}; v_i) &= \left(\frac{13}{14} - \frac{1}{4} \right) \cdot 0 = 0 \\ u_i(b_2, \boldsymbol{\beta}_{-i}; v_i) &= \left(\frac{13}{14} - \frac{1}{2} \right) (p_1/2 + (1 - p_1)/3) \\ u_i(b_3, \boldsymbol{\beta}_{-i}; v_i) &= \left(\frac{13}{14} - \frac{3}{4} \right) \cdot 1 \end{aligned}$$

Now, we have

– if $A[x] = 0$, i.e., $p_1 = 0$, then

$$u_i(b_3, \boldsymbol{\beta}_{-i}; v_i) - u_i(b_2, \boldsymbol{\beta}_{-i}; v_i) = \frac{5}{28} - \frac{3}{7} \cdot \frac{1}{3} = \frac{1}{28} > \varepsilon$$

and thus b_3 is the only ε -best-response, i.e., $\beta_{y'}(v_{y'})(b_3) = 1$. As proved above in the analysis of bidder y , this implies $A[y] = 1$.

– if $A[x] = 1$, i.e., $p_1 = 1$, then

$$u_i(b_2, \beta_{-i}; v_i) - u_i(b_3, \beta_{-i}; v_i) = \frac{3}{7} \cdot \frac{1}{2} - \frac{5}{28} = \frac{1}{28} > \varepsilon$$

and thus b_2 is the only ε -best-response, i.e., $\beta_{y'}(v_{y'})(b_2) = 1$. As proved above in the analysis of bidder y , this implies $A[y] = 0$.

As a result, the conditions of the NOT gate are indeed satisfied. \square

To complete the proof, note that by [Lemmas D.5 to D.7](#) all bidders $x \in X$ are valid, and thus, again by these lemmas, the constructed assignment A satisfies all the gates of the PURE-CIRCUIT instance (X, G) .

E Existence of Symmetric Equilibria in IID Auctions

This section is dedicated to proving [Theorem 5.2](#). En route, we establish two results that could be of independent interest, namely that (a) symmetric continuous first-price auctions (with discrete bidding space) always have pure symmetric equilibria (see [Theorem E.1](#)), and (b) the structure and approximability of (mixed) equilibria is preserved, in the limit, in discrete auctions (see [Lemma E.2](#)).

Theorem E.1. *Any first-price auction with continuous iid priors and discrete bids has a symmetric and monotone pure equilibrium.*

Proof. Our proof follows closely that of Athey [[Ath01](#)] for the existence of (pure) equilibria in continuous auctions for general (i.e., not necessarily symmetric) priors. We need to take care in handling correctly the symmetry of the underlying strategies, but this turns out to be straightforward.

Fix a CFPA \mathcal{A} with iid priors over $[0, 1]$ and discrete bidding space $B = \{b_1 \leq b_2 \leq \dots \leq b_m\}$. It is known that, in this continuous setting, any monotone strategy $\beta : [0, 1] \rightarrow B$ can be represented by a non-decreasing sequence of break points $\alpha = \alpha(\beta) = (\alpha_1, \alpha_2, \dots, \alpha_m)$ such that $\alpha_\ell = \sup\{v \in [0, 1] \mid \beta(v) \leq b_\ell\}$, for $\ell \in [m]$ (for a more detailed discussion, see [[FGHLP23](#), Sec. 2]). For simplicity, in the following we will slightly abuse notation, and sometimes use directly the jump-point representation $\alpha = \alpha(\beta)$ instead of a bidding strategy β .

Then, let us also denote the corresponding space of all feasible, no-overbidding (jump-point representation) strategies by

$$\Omega = \{\alpha \in [0, 1]^m \mid \alpha_\ell \leq \alpha_{\ell+1} \quad \forall \ell \in [m-1] \quad \wedge \quad b_\ell \leq \alpha_\ell \quad \forall \ell \in [m]\}.$$

Notice that the linearity of the inequalities defining Ω imply that it is a polytope, and thus Ω is convex and compact subspace of the Euclidean space \mathbb{R}^m . Also, it is easy to see that Ω is nonempty, since $(1, 1, \dots, 1) \in \Omega$ (corresponding to the strategy of always bidding $\min B = b_1$).

Next, given a symmetric and monotone strategy profile $\alpha = (\alpha, \alpha, \dots, \alpha)$, we define (the set of) its best-responses by

$$\text{BR}(\alpha) := \left\{ \alpha' \in \Omega \mid u_i(\alpha'(v_i), \alpha_{-i}; v_i) \geq \max_{b \in B} u_i(b, \alpha_{-i}; v_i) \quad \forall v_i \in V \right\}$$

Notice that, since auction \mathcal{A} has iid priors, the specific choice of the bidder i in the above definition is irrelevant. This gives rise to a correspondence $\text{BR} : \Omega \rightarrow 2^\Omega$ with the property that: a symmetric and monotone strategy profile, where all bidders play a strategy α , is an (exact) PBNE of \mathcal{A} , if and only if $\alpha \in \text{BR}(\alpha)$.

Using this property, we can establish the existence of a symmetric PBNE by means of Kakutani's fixed-point theorem (see, e.g., [[AB06](#), p. 583]). Towards that end, we need to establish the following properties:

- $\text{BR}(\alpha) \neq \emptyset$ for all $\alpha \in \Omega$.
- $\text{BR}(\alpha)$ is a convex set, for all $\alpha \in \Omega$.
- BR has a closed graph.

The above properties can be proved in an identical way to that of the existing results for general priors, and we refer the interested reader to [Ath01, Lemma 3] or [Kri09, p. 302]. \square

Lemma E.2 (Convergence Lemma). *Consider a DFPA \mathcal{A} with iid priors. Assume that \mathcal{A} has a symmetric monotone ε -approximate MBNE, for all $\varepsilon > 0$. Then, \mathcal{A} has an exact, symmetric and monotone, MBNE.*

Proof. Fix a DFPA \mathcal{A} with iid priors over a common value space V , and bidding space B . By the assumptions of our lemma, there exists a sequence of strategy profiles $\{\boldsymbol{\beta}_n = (\beta_n, \beta_n, \dots, \beta_n)\}$, $n \in \mathbb{N}$, where each bidder plays according to the same, monotone, strategy $\beta_n : V \rightarrow \Delta(B)$, such that $\boldsymbol{\beta}_n$ is a $\frac{1}{n}$ -approximate MBNE of \mathcal{A} for all positive integers n . For our proof, we will show that

- (a) there exists a subsequence $\{\bar{\beta}_n\}$ of $\{\beta_n\}$, such that $\bar{\beta}_n \rightarrow \beta^*$ (as $n \rightarrow \infty$), where β^* is a monotone strategy.
- (b) $\boldsymbol{\beta}^* = (\beta^*, \beta^*, \dots, \beta^*)$ is an *exact* MBNE.

Starting on (a), first observe that since we are at a discrete auction, there are finitely many different supports that a bidding strategy can have. Formally, since sets V, B are finite, the set $\{\times_{v \in V} \text{supp}(\beta(v))\}_{\beta \in \Delta(B)^V}$ is also finite. Therefore, there has to exist a subsequence $\{\tilde{\beta}_n\}$ of $\{\beta_n\}$ such that all its elements have exactly the same support.

Let $\tilde{p}_n = (\tilde{p}_n(v, b))_{v \in V, b \in B}$ denote the canonical representation of strategy $\tilde{\beta}_n$; that is, $\tilde{p}_n(v, b)$ is the probability of submitting bid b when having true value v , under strategy $\tilde{\beta}_n$ (see Section 2.1), for all $n \in \mathbb{N}$. Observe that all these probabilities must satisfy the feasibility constraints

$$\tilde{p}_n(v, b) \geq 0, \quad \sum_{b' \in B} \tilde{p}_n(v, b') = 1, \quad \text{and} \quad \tilde{p}_n(v, b)(v - b) \geq 0 \quad \forall v \in V, b \in B,$$

which define a polytope, and thus a compact subspace, of the Euclidean metric space $\mathbb{R}^{V \times B}$. Therefore, there has to exist a subsequence $\{\bar{p}_n\}$ of $\{\tilde{p}_n\}$ that converges to a feasible representation \bar{p}^* of bidding strategy $\beta^* \in \Delta(B)^V$.

The only thing remaining in order to establish point (a) is the monotonicity of β^* . Indeed, since all strategies of $\{\beta_n\}$ has the same support, it must be that *all* elements of the sequence $\{\bar{p}_n\}$ have exactly the same zero components. Therefore, due to compactness, the only way in which the supports of limiting strategy β^* may differ from the supports of all strategies in $\{\beta_n\}$ is if for a probability component it is $p^*(v, b) = 0$ but $\bar{p}_n(v, b) \neq 0$ (for all $n \in \mathbb{N}$). But then, this implies that the supports may only *shrink* in the limit; that is, $\text{supp}(\beta^*(v)) \subseteq \text{supp}(\bar{\beta}_n(v))$ for all $n \in \mathbb{N}$. So, given that $\bar{\beta}_n$'s are monotone, β^* must also be monotone; see Definition 3.

Moving now to point (b), for a symmetric strategy profile $\boldsymbol{\beta} = (\beta, \beta, \dots, \beta)$, let

$$Z(\boldsymbol{\beta}) := \max_{v \in V, b \in B} u_i(b, \boldsymbol{\beta}_{-i}; v_i) - u_i(\beta(v_i), \boldsymbol{\beta}_{-i}; v_i) \quad (43)$$

denote the maximum utility improvement that a bidder i can achieve by unilaterally deviating, from the bid $\beta(v_i)$ dictated by $\boldsymbol{\beta}$, to any other bid, while all other bidders play according to the symmetric profile $\boldsymbol{\beta} = (\beta, \beta, \dots, \beta)$, across all possible true values v_i . Note that, due to symmetry, the specific choice of the bidder i in the above definition (43) is irrelevant. Then, $\boldsymbol{\beta}$ is an ε -approximate MBNE of \mathcal{A} , if and only if $Z(\boldsymbol{\beta}) \leq \varepsilon$; see Definition 1. Obviously, $Z(\boldsymbol{\beta}) \geq 0$ for any strategy profile $\boldsymbol{\beta}$.

Observe that, as defined by (43), $Z : \Delta(B)^V \rightarrow [0, 1]$ is a continuous function: bidder utilities are continuous (see (2)) and (finitely many) max-operators preserve continuity. So,

$$Z(\beta^*) = \lim_{n \rightarrow \infty} Z(\bar{\beta}_n) \leq \lim_{n \rightarrow \infty} \frac{1}{n} = 0,$$

meaning that indeed β^* is an exact MBNE of \mathcal{A} . □

Now we have all the necessary machinery in order to prove [Theorem 5.2](#):

Proof of Theorem 5.2. Fix a DFPA \mathcal{A} with iid priors and choose an arbitrary $\varepsilon \in (0, 1)$. Then, by making use of our computational equivalence¹¹ between discrete and continuous auctions from [Section 4.2](#), we can construct a CFPA \mathcal{A}' with iid priors (and the same, discrete, bidding space as \mathcal{A}) such that, any ε -approximate symmetric (and monotone) PBNE of \mathcal{A}' can be translated back to a symmetric and monotone ε -approximate MBNE of \mathcal{A} .

Thus, given that such ε -approximate symmetric PBNE are guaranteed to exist, due to [Theorem E.1](#), we also get guaranteed existence of an ε -approximate, symmetric and monotone MBNE in our original discrete auction \mathcal{A} . Since we chose $\varepsilon > 0$ arbitrarily, our Convergence [Lemma E.2](#) gives us the desired existence of a monotone *exact* MBNE in \mathcal{A} . □

References

- [AB06] Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. 3rd. Springer-Verlag, 2006. DOI: [10.1007/3-540-29587-9](#) (cit. on p. 47).
- [AB24] Mete Şeref Ahunbay and Martin Bichler. “On the Uniqueness of Bayesian Coarse Correlated Equilibria in Standard First-Price and All-Pay Auctions”. In: *CoRR* abs/2401.01185 (2024). arXiv: [2401.01185](#) (cit. on pp. 6, 25).
- [ACLS22] Michael Albert, Vincent Conitzer, Giuseppe Lopomo, and Peter Stone. “Mechanism design for correlated valuations: Efficient methods for revenue maximization”. In: *Operations Research* 70.1 (2022), pp. 562–584. DOI: [10.1287/OPRE.2020.2092](#) (cit. on p. 4).
- [Ath01] Susan Athey. “Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information”. In: *Econometrica* 69.4 (July 2001), pp. 861–889. DOI: [10.1111/1468-0262.00223](#) (cit. on pp. 3, 8, 28, 47, 48).
- [Aum61] Robert J Aumann. *Mixed and behavior strategies in infinite extensive games*. Princeton University Princeton, 1961 (cit. on p. 11).
- [Aum74] Robert J. Aumann. “Subjectivity and correlation in randomized strategies”. In: *Journal of Mathematical Economics* 1.1 (1974), pp. 67–96. ISSN: 0304-4068. DOI: [10.1016/0304-4068\(74\)90037-8](#) (cit. on pp. 6, 25).
- [Aum87] Robert J Aumann. “Correlated equilibrium as an expression of Bayesian rationality”. In: *Econometrica: Journal of the Econometric Society* (1987), pp. 1–18. DOI: [10.2307/1911154](#) (cit. on pp. 6, 25).
- [BDMMZ21] Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. “Robust Auction Design in the Auto-bidding World”. In: *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS)*. 2021 (cit. on p. 3).

¹¹As a matter of fact, in this proof we make use of the reduction from discrete to continuous auctions by deploying [Lemma 4.2](#) with $\varepsilon \leftarrow 0$ and $\delta \leftarrow \varepsilon$ in its statement.

- [BG97] Pierpaolo Battigalli and Danilo Guitoli. “Conjectural Equilibria and Rationalizability in a Game with Incomplete Information”. In: *Decisions, Games and Markets*. Ed. by Pierpaolo Battigalli, Aldo Montesano, and Fausto Panunzi. Springer, 1997, pp. 97–124. DOI: [10.1007/978-1-4615-6337-2_4](https://doi.org/10.1007/978-1-4615-6337-2_4) (cit. on p. 4).
- [BGM92] Pierpaolo Battigalli, Mario Gilli, and M. Cristina Molinari. “Learning and Convergence to Equilibrium in Repeated Strategic Interactions: An Introductory Survey”. In: *Ricerche Economiche* 46 (1992), pp. 335–378 (cit. on p. 4).
- [BM13] Dirk Bergemann and Stephen Morris. “Robust predictions in games with incomplete information”. In: *Econometrica* 81.4 (2013), pp. 1251–1308. DOI: [10.2139/ssrn.2163606](https://doi.org/10.2139/ssrn.2163606) (cit. on pp. 6, 25).
- [BM16] Dirk Bergemann and Stephen Morris. “Bayes correlated equilibrium and the comparison of information structures in games”. In: *Theoretical Economics* 11.2 (2016), pp. 487–522 (cit. on pp. 6, 25).
- [Cas67] Ralph Cassady. *Auctions and auctioneering*. University of California Press, 1967 (cit. on p. 3).
- [CDT09] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. “Settling the Complexity of Computing Two-player Nash Equilibria”. In: *Journal of the ACM* 56.3 (May 2009), 14:1–14:57. DOI: [10.1145/1516512.1516516](https://doi.org/10.1145/1516512.1516516) (cit. on p. 3).
- [CDW12] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. “An algorithmic characterization of multi-dimensional mechanisms”. In: *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC)*. 2012, pp. 459–478. DOI: [10.1145/2213977.2214021](https://doi.org/10.1145/2213977.2214021) (cit. on p. 4).
- [CH13] Shuchi Chawla and Jason D. Hartline. “Auctions with unique equilibria”. In: *Proceedings of the 14th ACM conference on Electronic Commerce (EC)*. 2013, pp. 181–196. DOI: [10.1145/2492002.2483188](https://doi.org/10.1145/2492002.2483188) (cit. on p. 19).
- [Chw89] Michael Suk-Young Chwe. “The Discrete Bid First Auction”. In: *Economics Letters* 31.4 (Dec. 1989), pp. 303–306. DOI: [10.1016/0165-1765\(89\)90019-0](https://doi.org/10.1016/0165-1765(89)90019-0) (cit. on pp. 3, 19).
- [CKSS22] Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E Stier-Moses. “Multiplicative pacing equilibria in auction markets”. In: *Operations Research* 70.2 (2022), pp. 963–989. DOI: [10.1287/opre.2021.2167](https://doi.org/10.1287/opre.2021.2167) (cit. on p. 3).
- [CP14] Yang Cai and Christos Papadimitriou. “Simultaneous Bayesian auctions and computational complexity”. In: *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*. June 2014, pp. 895–910. DOI: [10.1145/2600057.2602877](https://doi.org/10.1145/2600057.2602877) (cit. on p. 6).
- [CP23] Xi Chen and Binghui Peng. “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)*. 2023, pp. 698–709. DOI: [10.1145/3564246.3585195](https://doi.org/10.1145/3564246.3585195) (cit. on pp. 3, 5, 6, 8, 11, 12, 18, 28).
- [CRS82] James C Cox, Bruce Roberson, and Vernon L Smith. “Theory and behavior of single object auctions”. In: *Research in experimental economics* 2.1 (1982), pp. 1–43. DOI: [10.1017/CB09780511528354.031](https://doi.org/10.1017/CB09780511528354.031) (cit. on p. 4).
- [CS08] Vincent Conitzer and Tuomas Sandholm. “New complexity results about Nash equilibria”. In: *Games and Economic Behavior* 63.2 (2008), pp. 621–641. DOI: [10.1016/j.geb.2008.02.015](https://doi.org/10.1016/j.geb.2008.02.015) (cit. on p. 6).

- [CST80] Vicki M Coppinger, Vernon L Smith, and Jon A Titus. “Incentives and behavior in english, dutch and sealed-bid auctions”. In: *Economic inquiry* 18.1 (1980), pp. 1–22. DOI: [10.1111/J.1465-7295.1980.TB00556.X](https://doi.org/10.1111/J.1465-7295.1980.TB00556.X) (cit. on p. 4).
- [CSW83] James C Cox, Vernon L Smith, and James M Walker. “Tests of a heterogeneous bidders theory of first price auctions”. In: *Economics Letters* 12.3-4 (1983), pp. 207–212. DOI: [10.1016/0165-1765\(83\)90039-3](https://doi.org/10.1016/0165-1765(83)90039-3) (cit. on p. 4).
- [CSW88] James C Cox, Vernon L Smith, and James M Walker. “Theory and individual behavior of first-price auctions”. In: *Journal of Risk and uncertainty* 1 (1988), pp. 61–99. DOI: [10.1007/BF00055565](https://doi.org/10.1007/BF00055565) (cit. on p. 4).
- [DFHM22] Argyrios Deligkas, John Fearnley, Alexandros Hollender, and Themistoklis Melissourgos. “Pure-circuit: Strong inapproximability for PPAD”. In: *Proceedings of the 63rd IEEE Symposium on Foundations of Computer Science (FOCS)*. 2022, pp. 159–170. DOI: [10.1109/FOCS54457.2022.00022](https://doi.org/10.1109/FOCS54457.2022.00022) (cit. on pp. 5, 18, 41).
- [DFP06] Constantinos Daskalakis, Alex Fabrikant, and Christos H. Papadimitriou. “The Game World Is Flat: The Complexity of Nash Equilibria in Succinct Games”. In: *Proceedings of the 33rd International Colloquium on Automata, Languages, and Programming (ICALP)*. Ed. by Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener. 2006, pp. 513–524. DOI: [10.1007/11786986_45](https://doi.org/10.1007/11786986_45) (cit. on p. 5).
- [DGP09] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. “The complexity of computing a Nash equilibrium”. In: *SIAM Journal on Computing* 39.1 (2009), pp. 195–259. DOI: [10.1137/070699652](https://doi.org/10.1137/070699652) (cit. on p. 3).
- [Dig19] Digiday.com. *Digiday Research: Header bidding and first-price auctions boost publisher revenues*. Accessed: 2019-01-14. 2019. URL: https://digiday.com/media/digiday-research-header-bidding-and-first-price-auctions-boost-publisher-revenues/?utm_medium=email&utm_campaign=digidaydis&utm_source=uk&utm_content=190114 (cit. on p. 3).
- [DRS21] Stylianos Despotakis, Ramamoorthi Ravi, and Amin Sayedi. “First-price auctions in online display advertising”. In: *Journal of Marketing Research* 58.5 (2021), pp. 888–907. DOI: [10.1177/00222437211030201](https://doi.org/10.1177/00222437211030201) (cit. on p. 3).
- [EMS09] Guillaume Escamocher, Peter Bro Miltersen, and Rocio Santillan R. “Existence and Computation of Equilibria of First-price Auctions with Integral Valuations and Bids”. In: *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2009, pp. 1227–1228. URL: <https://dl.acm.org/doi/10.5555/1558109.1558225> (cit. on pp. 4–6, 8, 10).
- [FGHLP23] Aris Filos-Ratsikas, Yiannis Giannakopoulos, Alexandros Hollender, Philip Lazos, and Diogo Poças. “On the Complexity of Equilibrium Computation in First-Price Auctions”. In: *SIAM Journal on Computing* 52.1 (2023), pp. 80–131. DOI: [10.1137/21M1435823](https://doi.org/10.1137/21M1435823) (cit. on pp. 3–6, 8, 9, 11, 12, 17, 18, 47).
- [FHHK14] Hu Fu, Nima Haghpanah, Jason Hartline, and Robert Kleinberg. “Optimal auctions for correlated buyers with sampling”. In: *Proceedings of the fifteenth ACM conference on Economics and computation*. 2014, pp. 23–36. DOI: [10.1145/2600057.2602895](https://doi.org/10.1145/2600057.2602895) (cit. on p. 4).
- [FL86] Drew Fudenberg and David Levine. “Limit Games and Limit Equilibria”. In: *Journal of Economic Theory* 38.2 (Apr. 1986), pp. 261–279. DOI: [10.1016/0022-0531\(86\)90118-3](https://doi.org/10.1016/0022-0531(86)90118-3) (cit. on p. 4).

- [For06] Françoise Forges. “Correlated equilibrium in games with incomplete information revisited”. In: *Theory and Decision* 61 (Dec. 2006), pp. 329–344. DOI: [10.1007/s11238-006-9005-3](https://doi.org/10.1007/s11238-006-9005-3) (cit. on pp. 6, 25, 27).
- [For23] Françoise Forges. “Correlated equilibrium in games with incomplete information”. In: *Revue Économique* 74 (2023), pp. 529–539. DOI: [10.3917/reco.744.0529](https://doi.org/10.3917/reco.744.0529) (cit. on pp. 6, 25, 27).
- [For93] Françoise Forges. “Five legitimate definitions of correlated equilibrium in games with incomplete information”. In: *Theory and Decision* 35 (1993), pp. 277–310. URL: <https://api.semanticscholar.org/CorpusID:119900539> (cit. on pp. 6, 25, 27).
- [FW16] Rafael Frongillo and Jens Witkowski. “A Geometric Method to Construct Minimal Peer Prediction Mechanisms”. In: *Proceedings of the 30th AAAI Conference on Artificial Intelligence*. 2016, pp. 502–508. DOI: [10.5555/3015812.3015888](https://doi.org/10.5555/3015812.3015888) (cit. on p. 4).
- [GGM07] Georg Gottlob, Gianluigi Greco, and Toni Mancini. “Complexity of Pure Equilibria in Bayesian Games”. In: *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*. 2007, pp. 1294–1299. DOI: [10.5555/1625275.1625485](https://doi.org/10.5555/1625275.1625485) (cit. on p. 6).
- [GJ79] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1979 (cit. on pp. 10, 32).
- [GLS67] James H. Griesmer, Richard E. Levitan, and Martin Shubik. “Toward a Study of Bidding Processes Part IV – Games With Unknown Costs”. In: *Naval Research Logistics* 14.4 (1967), pp. 415–433. DOI: [10.1002/nav.3800140402](https://doi.org/10.1002/nav.3800140402) (cit. on p. 3).
- [Gol11] Paul W. Goldberg. “A survey of PPAD-completeness for computing Nash equilibria”. In: *Surveys in Combinatorics 2011*. Ed. by Robin Chapman. London Mathematical Society Lecture Note Series. Cambridge University Press, 2011, pp. 51–82. DOI: [10.1017/CBO9781139004114.003](https://doi.org/10.1017/CBO9781139004114.003) (cit. on p. 4).
- [GV88] D. Yu. Grigor’ev and N.N. Vorobjov. “Solving systems of polynomial inequalities in subexponential time”. In: *Journal of Symbolic Computation* 5.1-2 (1988), pp. 37–64. DOI: [10.1016/s0747-7171\(88\)80005-1](https://doi.org/10.1016/s0747-7171(88)80005-1) (cit. on pp. 6, 20, 22).
- [Hah73] Frank Hahn. *On the Notion of Equilibrium in Economics: An Inaugural Lecture by F.H. Hahn*. Cambridge University Press, 1973 (cit. on p. 4).
- [Har67] John C. Harsanyi. “Games With Incomplete Information Played by “Bayesian” Players, I–III: Part I. The Basic Model”. In: *Management Science* 14.3 (1967), pp. 159–182. DOI: [10.1287/mnsc.1040.0270](https://doi.org/10.1287/mnsc.1040.0270) (cit. on pp. 3, 4, 25).
- [Har89] Glenn W Harrison. “Theory and misbehavior of first-price auctions”. In: *The American Economic Review* (1989), pp. 749–762 (cit. on p. 4).
- [HST15] Jason Hartline, Vasilis Syrgkanis, and Eva Tardos. “No-regret learning in Bayesian games”. In: *Advances in Neural Information Processing Systems* 28 (2015) (cit. on pp. 6, 25).
- [JL15] Albert Xin Jiang and Kevin Leyton-Brown. “Polynomial-time computation of exact correlated equilibrium in compact games”. In: *Games and Economic Behavior* 91 (2015), pp. 347–359. ISSN: 0899-8256. DOI: [10.1016/j.geb.2013.02.002](https://doi.org/10.1016/j.geb.2013.02.002) (cit. on pp. 25, 26).
- [JR01] Geoffrey A. Jehle and Philip J. Reny. *Advanced Microeconomic Theory*. Financial Times/Prentice Hall, 2001 (cit. on pp. 4, 6–8, 11, 25).
- [Kak41] Shizuo Kakutani. “A generalization of Brouwer’s fixed point theorem”. In: *Duke Mathematical Journal* 8.3 (Sept. 1941). DOI: [10.1215/s0012-7094-41-00838-4](https://doi.org/10.1215/s0012-7094-41-00838-4) (cit. on p. 19).

- [KBP18] Çağlı Koçyiğit, Halil I Bayrak, and Mustafa Ç Pınar. “Robust auction design under multiple priors by linear and integer programming”. In: *Annals of Operations Research* 260.1-2 (2018), pp. 233–253. DOI: [10.1007/s10479-017-2416-4](https://doi.org/10.1007/s10479-017-2416-4) (cit. on p. 4).
- [KL93] Ehud Kalai and Ehud Lehrer. “Rational Learning Leads to Nash Equilibrium”. In: *Econometrica* 61.5 (1993), pp. 1019–1045. DOI: [10.2307/2951492](https://doi.org/10.2307/2951492) (cit. on p. 4).
- [Kri09] Vijay Krishna. *Auction Theory*. 2nd. Academic Press, 2009 (cit. on pp. 19, 28, 48).
- [LB17] Brendan Lucier and Allan Borodin. “Equilibria of greedy combinatorial auctions”. In: *SIAM Journal on Computing* 46.2 (2017), pp. 620–660. DOI: [10.1137/15M1048720](https://doi.org/10.1137/15M1048720) (cit. on p. 7).
- [Leb06] Bernard Lebrun. “Uniqueness of the equilibrium in first-price auctions”. In: *Games and Economic Behavior* 55.1 (Apr. 2006), pp. 131–151. DOI: [10.1016/j.geb.2005.01.006](https://doi.org/10.1016/j.geb.2005.01.006) (cit. on pp. 3, 8).
- [Leb96] Bernard Lebrun. “Existence of an Equilibrium in First Price Auctions”. In: *Economic Theory* 7 (1996), pp. 421–443. DOI: [10.1007/BF01213659](https://doi.org/10.1007/BF01213659) (cit. on pp. 3, 28).
- [Leb99] Bernard Lebrun. “First Price Auctions in the Asymmetric N Bidder Case”. In: *International Economic Review* 40.1 (1999), pp. 125–142. DOI: [10.1111/1468-2354.00008](https://doi.org/10.1111/1468-2354.00008) (cit. on p. 3).
- [LP00] Alessandro Lizzeri and Nicola Persico. “Uniqueness and Existence of Equilibrium in Auctions with a Reserve Price”. In: *Games and Economic Behavior* 30.1 (Jan. 2000), pp. 83–114. DOI: [10.1006/game.1998.0704](https://doi.org/10.1006/game.1998.0704) (cit. on p. 3).
- [MNT09] Peter Bro Miltersen, Jesper Buus Nielsen, and Nikos Triandopoulos. “Privacy-enhancing auctions using rational cryptography”. In: *Annual International Cryptology Conference*. Springer, 2009, pp. 541–558. DOI: [10.1145/1807406.1807496](https://doi.org/10.1145/1807406.1807496) (cit. on p. 4).
- [MP91] Nimrod Megiddo and Christos H. Papadimitriou. “On total functions, existence theorems and computational complexity”. In: *Theoretical Computer Science* 81.2 (1991), pp. 317–324. DOI: [10.1016/0304-3975\(91\)90200-1](https://doi.org/10.1016/0304-3975(91)90200-1) (cit. on p. 17).
- [MR00] Eric Maskin and John Riley. “Equilibrium in Sealed High Bid Auctions”. In: *The Review of Economic Studies* 67.3 (2000), pp. 439–454. DOI: [10.1111/1467-937X.00138](https://doi.org/10.1111/1467-937X.00138) (cit. on pp. 3, 6, 8, 28).
- [MR03] Eric Maskin and John Riley. “Uniqueness of equilibrium in sealed high-bid auctions”. In: *Games and Economic Behavior* 45.2 (2003), pp. 395–409. DOI: [10.1016/S0899-8256\(03\)00150-7](https://doi.org/10.1016/S0899-8256(03)00150-7) (cit. on pp. 3, 8).
- [MR85] Eric S. Maskin and John G. Riley. “Auction theory with private values”. In: *The American Economic Review* 75.2 (1985), pp. 150–155. URL: <https://www.jstor.org/stable/1805587> (cit. on pp. 3–5, 8, 28).
- [MW85] Paul R. Milgrom and Robert J. Weber. “Distributional strategies for games with incomplete information”. In: *Mathematics of Operations Research* 10.4 (1985), pp. 619–632. DOI: [10.1287/moor.10.4.619](https://doi.org/10.1287/moor.10.4.619) (cit. on pp. 7, 11).
- [Mye97] Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1997 (cit. on pp. 4, 6, 7, 11, 25, 28).
- [Nas51] John Nash. “Non-cooperative games”. In: *Annals of mathematics* (1951), pp. 286–295 (cit. on p. 11).
- [Pap94] Christos H. Papadimitriou. “On the complexity of the parity argument and other inefficient proofs of existence”. In: *Journal of Computer and System Sciences* 48.3 (1994), pp. 498–532. DOI: [10.1016/s0022-0000\(05\)80063-7](https://doi.org/10.1016/s0022-0000(05)80063-7) (cit. on p. 3).

- [Plu92] Michael Plum. “Characterization and Computation of Nash-Equilibria for Auctions with Incomplete Information”. In: *International Journal of Game Theory* 20.4 (Dec. 1992), pp. 393–418. DOI: [10.1007/bf01271133](https://doi.org/10.1007/bf01271133) (cit. on p. 3).
- [PP15] Christos Papadimitriou and George Pierrakos. “Optimal deterministic auctions with correlated priors”. In: *Games and Economic Behavior* 92 (2015), pp. 430–454. DOI: [10.1016/j.geb.2013.08.009](https://doi.org/10.1016/j.geb.2013.08.009) (cit. on p. 4).
- [PR08] Christos H. Papadimitriou and Tim Roughgarden. “Computing Correlated Equilibria in Multi-Player Games”. In: *Journal of the ACM* 55.3 (Aug. 2008), pp. 141–1429. DOI: [10.1145/1379759.1379762](https://doi.org/10.1145/1379759.1379762) (cit. on pp. 6, 25, 26).
- [PST20] Renato Paes Leme, Balasubramanian Sivan, and Yifeng Teng. “Why Do Competitive Markets Converge to First-Price Auctions?” In: *Proceedings of The World Wide Web Conference (WWW)*. 2020, pp. 596–605. DOI: [10.1145/3366423.3380142](https://doi.org/10.1145/3366423.3380142) (cit. on p. 3).
- [RS81] John G. Riley and William F. Samuelson. “Optimal Auctions”. In: *The American Economic Review* 71.3 (1981), pp. 381–392. URL: <https://www.jstor.org/stable/1802786> (cit. on p. 19).
- [RW94] Ariel Rubinstein and Asher Wolinsky. “Rationalizable Conjectural Equilibrium: Between Nash and Rationalizability”. In: *Games and Economic Behavior* 6.2 (Mar. 1994), pp. 299–311. DOI: [10.1006/game.1994.1016](https://doi.org/10.1006/game.1994.1016) (cit. on p. 4).
- [Sta24] Statista.com. *Annual revenue of Alphabet from 2017 to 2023, by segment*. Accessed: 2024-02-12. 2024. URL: <https://www.statista.com/statistics/633651/alphabet-annual-global-revenue-by-segment/> (cit. on p. 3).
- [Vic61] William Vickrey. “Counterspeculation, Auctions and Competitive Sealed Tenders”. In: *Journal of Finance* 16.1 (Mar. 1961), pp. 8–37. DOI: [10.1111/j.1540-6261.1961.tb02789.x](https://doi.org/10.1111/j.1540-6261.1961.tb02789.x) (cit. on p. 3).
- [WP12] Jens Witkowski and David C. Parkes. “Peer prediction without a common prior”. In: *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*. 2012, pp. 964–981. DOI: [10.1145/2229012.2229085](https://doi.org/10.1145/2229012.2229085) (cit. on p. 4).
- [WSZ20] Zihe Wang, Weiran Shen, and Song Zuo. “Bayesian Nash Equilibrium in First-Price Auction with Discrete Value Distributions”. In: *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2020, pp. 1458–1466. URL: <https://dl.acm.org/doi/abs/10.5555/3398761.3398929> (cit. on pp. 6, 8).