The Price of Stability for First Price Auction*

Yaonan Jin† Pinyan Lu‡

Abstract

This paper establishes the Price of Stability (PoS) for First Price Auctions, for all equilibrium concepts that have been studied in the literature: Bayesian Nash Equilibrium $\subsetneq$ Bayesian Correlated Equilibrium $\subsetneq$ Bayesian Coarse Correlated Equilibrium.

- Bayesian Nash Equilibrium: For independent valuations, the tight PoS is $1 - 1/e^2 \approx 0.8647$, matching the counterpart Price of Anarchy (PoA) bound [JL22]. For correlated valuations, the tight PoS is $1 - 1/e \approx 0.6321$, matching the counterpart PoA bound [ST13, Syr14].

  This result indicates that, in the worst cases, efficiency degradation depends not on different selections among Bayesian Nash Equilibria.

- Bayesian (Coarse) Correlated Equilibrium: For independent or correlated valuations, the tight PoS is always $1 = 100\%$, i.e., no efficiency degradation.

  This result indicates that First Price Auctions can be fully efficient when we allow the more general equilibrium concepts.

1 Introduction.

It is well-known in game theory that a multi-agent system might be in suboptimal states due to selfish behavior of the agents. Auctions are an important genre of such systems. In a single-item auction, each bidder $i \in [n]$ independently draws her value from a distribution $v_i \sim V_i$ but does not know others' values $\boldsymbol{v}_{-i} = (v_k)_{k \neq i}$. Then, each bidder i submits a (possibly random) bid $b_i = s_i(v_i)$ based on her value v_i and strategy s_i. The auction rule determines the winner and how much the bidders need to pay. Each bidder i has a quasi-linear utility function $u_i(v_i,\, b_i) = v_i \cdot \mathrm{x}_i(b_i) - \rho_i(b_i)$, where the winning probability $\mathrm{x}_i(b_i)$ and the expected payment $\rho_i(b_i)$ are taken over the randomness of other bidders' values and strategies, as well as the inherent randomness of the auction. Like other game-theoretical systems, we can define the *equilibria* of an auction.

Definition 1.1. (Equilibria) *A strategy profile $\boldsymbol{s} = \{s_i\}_{i \in [n]}$ is a Bayesian Nash Equilibrium for an auction $\mathcal{A}$ when: For each bidder $i \in [n]$ and any possible value $v \in \mathrm{supp}(V_i)$, the considered strategy $s_i(v)$ is optimal, namely $\mathbf{E}_{s_i}\left[u_i(v,\, s_i(v))\right] \geq u_i(v,\, b)$ for any deviation bid $b \geq 0$. Denote by $\mathbb{BNE}(\boldsymbol{V})$ the space of Bayesian Nash Equilibria of an instance $\boldsymbol{V} = \{V_i\}_{i \in [n]}$.*

Auctions are widely employed to allocate recourse in a competitive environment. Thus *efficiency* is a central property of an auction. Given an auction $\mathcal{A}$, the social welfare from an instance $\boldsymbol{V}$ at a specific equilibrium $\boldsymbol{s}$, denoted by $\mathcal{A}(\boldsymbol{V},\, \boldsymbol{s})$, is the expectation of the winner's value. Ideally, we would like to allocate the item always to the bidder who values it the most; in expectation, this gives the optimal social welfare $\mathsf{OPT}(\boldsymbol{V})$.

For many auctions, the auction social welfare $\mathcal{A}(\boldsymbol{V},\, \boldsymbol{s})$ in general is strictly below the optimal social welfare $\mathsf{OPT}(\boldsymbol{V})$. To measure the (in)efficiency of an auction $\mathcal{A}$, we can define its Price of Anarchy [KP99] as the *worst-case ratio* between the two social welfares.

Definition 1.2. (Price of Anarchy) *The Price of Anarchy of an auction $\mathcal{A}$ is given by*

$$\mathsf{PoA} \;:=\; \inf_{\boldsymbol{V}} \inf_{\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V})} \left\{ \frac{\mathcal{A}(\boldsymbol{V},\, \boldsymbol{s})}{\mathsf{OPT}(\boldsymbol{V})} \right\}.$$

*The full version of the paper can be accessed at https://arxiv.org/abs/2207.04455

†Columbia University. Email: yj2552@columbia.edu

‡ITCS, Shanghai University of Finance and Economics. Email: lu.pinyan@mail.shufe.edu.cn

We consider the (in)efficiency problem for the *first-price auction*, one of the most widely-used auctions. The rule of the first-price auction is very simple: The bidder with the highest bid wins and pays her bid. Simple as the rule is, it is well-known that the equilibria can be very complicated. E.g. ([Vic61]), suppose that there are only two bidders, Alice has a [0, 1]-uniform random value v_1 and Bob has a [0, 2]-uniform random value v_2, then the unique Bayesian Nash Equilibrium takes the form of $s_1(v_1) = \frac{4}{3v_1}\big(1 - \sqrt{1 - \frac{3}{4}v_1^2}\big)$ and $s_2(v_2) = \frac{4}{3v_2}\big(\sqrt{1 + \frac{3}{4}v_2^2} - 1\big)$.

For the PoA in the first-price auction, Syrgkanis and Tardos [ST13] obtained the first nontrivial lower bound of $1 - 1/e \approx 0.6321$. Later, Hoy, Taggart, and Wang [HTW18] gave an improved lower bound of ≈ 0.7430. In a recent work by the authors [JL22], the tight bound of $1 - 1/e^2 \approx 0.8647$ was finally derived. That is a complete and insightful characterization. However, there are still a few remaining issues about the efficiency of the first-price auction, which we will discuss and address in this paper.

First, it is well-known that certain instances may have *no* equilibrium. For those instances, the tight PoA bound by [JL22] does not imply anything about the efficiency of the first-price auction. Given this, the natural question is, to what extent can we generalize the tight PoA results?

Second, it is also well-known that certain instances may have *multiple* or even *infinite* equilibria. For those instances, Price of Anarchy may be too *pessimistic* a measure since it concentrates just on the worst-case equilibria. Especially, the worst-case instance by [JL22] for the tight $\mathsf{PoA} = 1 - 1/e^2$ does have other more efficient or even fully efficient equilibria. Towards an *optimistic* measure of (in)efficiency, we shall consider another widely studied concept called Price of Stability [ADK+08], which is targeted at the *best-case* equilibria (instead of the *worst-case* equilibria as for PoA).

Definition 1.3. (Price of Stability) *The Price of Stability of an auction $\mathcal{A}$ is given by*

$$\mathsf{PoS} \;:=\; \inf_{\boldsymbol{V}} \sup_{\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V})} \left\{ \frac{\mathcal{A}(\boldsymbol{V},\, \boldsymbol{s})}{\mathsf{OPT}(\boldsymbol{V})} \right\}.$$

By definition, the tight PoS must be lower bounded by the tight PoA. Especially, for the first-price auction, we have $1 - 1/e^2 \le \mathsf{PoS} \le 1$.

Third, there are other modelings of the (in)efficiency problem. I.e., the above canonical setting assumes (bidder-wise) *independent* valuations $\boldsymbol{V}$ and strategies $\boldsymbol{s}$. Instead, one can consider *correlated valuations*, which is quite common in real life. Also, one can consider *correlated strategies*, for which the counterpart solution concepts are (i) Bayesian Correlated Equilibrium and (ii) Bayesian Coarse Correlated Equilibrium;[1] see Appendix A for the formal definitions. In total, we have two valuation classes and three equilibrium concepts, thus $2 \times 3 = 6$ meaningful settings. In *each* setting, the PoA and the PoS are both of fundamental interest.

The previous literature studies more on PoA and the tight bounds have been obtained in most settings; see Section 1.2 for a detailed review. In contrast, the tight PoS bounds remain open in all settings. (Maybe this is because, in each setting, the PoS as the solution to a *minimax optimization* problem shall be more challenging than the PoA as the solution to a *minimization* problem.) And understanding those PoS bounds is the main focus of our work.

Besides the concrete bounds, it is also interesting to know in which settings the PoS coincides with the PoA. Namely, if they are equal $\mathsf{PoA} = \mathsf{PoS}$, then this bound is a better characterization of the efficiency since it is "robust" against different equilibria.

1.1 Our results.

In this work, we will address each of the three issues mentioned above.

For the (possible) non-existence of Bayesian Nash Equilibria in the first-price auction, we show that this is just a consequence of "the underlying *tie-breaking rule* of the auction is incompatible with the considered value distribution $\boldsymbol{V}$". As a remedy, we prove that, for any $\delta > 0$, there always exists a δ-approximate Bayesian Nash Equilibrium that (i) makes *any* tie-breaking rule compatible with the considered value distribution $\boldsymbol{V}$, and (ii) the resulting auction social welfare is at least a $1 - 1/e^2$ fraction of the optimal social welfare. This indicates that the PoA characterization of the first-price auction is robust and universal. See Section 3 for more details.

[1] It is well-known (see [Rou15]) that Bayesian Nash Equilibrium is more special than Bayesian Correlated Equilibrium, which then is more special than Bayesian Coarse Correlated Equilibrium.

Different definitions of Bayesian (Coarse) Correlated Equilibria are considered in the literature, such as those by [CKK+15] vs. those by [ST13]; for a thorough discussion, the reader can refer to [Syr14, Chapter 3.3.1]. This paper will use the definitions by [CKK+15].

The main result of our work is the tight PoS bounds in all settings, summarized as follows.

	Independent Valuations	Correlated Valuations
BNE	PoS $= 1 - 1/e^2$ [Theorem 4.1]	PoS $= 1 - 1/e$ [Theorem 5.1]
BCE	PoS $= 1$ [Theorem A.1]	
BCCE		

Interestingly, in the settings of Bayesian Nash Equilibrium for either independent or correlated valuations, the tight PoS bounds coincide with the PoA counterparts (see Table 1). This would be an easy corollary if the known PoA-worst instances, due to [JL22] and [Syr14] respectively, each have unique equilibria. Unfortunately, this is not the case for the either instance. Even worse, the either instance has *fully efficient* equilibria, so the PoS bound thereof is 1.

Towards the tight PoS bounds $= 1 - 1/e^2$ or $1 - 1/e$, we shall modify the (original) PoA-worst instances from [JL22] and [Syr14]. For each modified instance, we first show and verify a particular equilibrium, named by the *focal* equilibrium $\boldsymbol{s}^*$, that is adjusted from the *worst-case* equilibrium for the original instance. More importantly, unlike the original instance, the modification eliminates other *more efficient* equilibria, left only with the focal equilibrium $\boldsymbol{s}^*$. Namely, we prove that the focal equilibrium $\boldsymbol{s}^*$ is the unique Bayesian Nash Equilibrium of the modified instance. Furthermore, the modification can be small enough in magnitude, such that the modified auction/optimal social welfares are arbitrarily close to the original counterparts. As a combination, we obtain the identity PoS $=$ PoA $= 1 - 1/e^2$ or $1 - 1/e$ in the either setting. See Sections 4 and 5 for more details.

That PoA and PoS have the same tight bounds is conceptually important – Such a PoA $=$ PoS tight bound "truly" captures the worst-case efficiency of Bayesian Nash Equilibria in the first-price action, despite the variety of equilibria and the selection among equilibria.

For Bayesian Correlated Equilibrium and/or Bayesian Coarse Correlated Equilibrium, we show that there always exist *fully efficient* equilibria. So, in those settings, whether independent or correlated valuations, we always have PoS $= 1$. (Notice that regarding a more general equilibrium concept, the PoS becomes larger, while the PoA becomes smaller.) See Appendix A for more details.

1.2 Related works. The first-price auction and its efficiency, motivated by its overwhelming prevalence in real business, are centerpiece of modern auction theory. This study dates back to Vickrey's seminal paper [Vic61] and has cultivated a rich literature [SZ90, Plu92, Leb96, Leb99, MR00a, MR00b, JSSZ02, MR03, Leb06, HKMN11, KZ12, CH13, and the references therein]. However, those works are restricted to special cases – The equilibria in the first-price auction are notoriously complicated; thus in general, classical economic analysis suffers from certain obstacles. From a computational perspective, there also is evidence for why the equilibria are elusive [CP14, FGH+21].

Over the last two decades, works from computer science bring a fresh viewpoint, *approximation guarantees* at the worst-/best-case equilibria, thus coining the concepts "Price of Anarchy/Stability" [KP99, ADK+08]. Regarding the first-price auction and Bayesian Nash Equilibria, the state-of-the-art results are summarized in Table 1.

	Deterministic Valuations		Independent Valuations		Correlated Valuations	
PoA	TB $= 1$	folklore	TB $= 1 - \frac{1}{e^2}$	[JL22]	TB $= 1 - \frac{1}{e}$	[ST13, Syr14]
PoS				Theorem 4.1		Theorem 5.1

Table 1: Tight PoA/PoS bounds regarding Bayesian Nash Equilibria.

Technically, the most prevalent tool for studying Price of Anarchy in auctions is the smoothness framework proposed by Roughgarden [Rou15] and then developed by Syrgkanis and Tardos [ST13]. This framework enables the tight bound $= 1 - 1/e$ in most settings, but has inherent bottlenecks in the canonical setting, namely Bayesian Nash Equilibrium for independent valuations. To mitigate those issues, Hoy, Taggart, and Wang [HTW18] combined additional techniques into the smoothness framework, hence an improved lower bound of ≈ 0.7430.

Until very recently, through a completely new framework, the authors [JL22] finally derived the tight bound $= 1 - 1/e^2 \approx 0.8647$.

The above discussions all concern efficiency guarantees. Another interesting and relevant topic is revenue guarantees in the first-price auction. Hartline, Hoy, and Taggart [HHT14] showed that, when the auctioneer sets *bidder-personalized* reserves in the first-price auction, the worst-case equilibria achieve a $\geq \frac{1}{2}(1 - 1/e) \approx 31.61\%$ approximation to optimal revenues. As an implication of the later works [AHN+19, JLQ+19], a better revenue guarantee $\gtrsim \frac{1}{2.6202} \approx 38.17\%$ holds even when the auctioneer sets *bidder-anonymous* reserves. It would be interesting to capture the revenue-PoA and revenue-PoS for the first-price auction with (optimal) personalized/anonymous reserves.

2 Notation and Preliminaries.

This section presents a bunch of structural results from the literature, especially [JL22], which lay the foundation of our paper. (More structural results will be presented in the later sections, when they are needed for our discussions.)

In a single-item auction, the bidders $[n] = \{1, 2, \ldots, n\}$ submit *non-negative* bids $\boldsymbol{b} = (b_i)_{i \in [n]}$ to the auctioneer. First Price Auction is a family of auctions $\mathcal{A} = (\mathrm{x}, \boldsymbol{\rho})$ that all obey the first-price allocation/payment principles.

- **first-price allocation:** Let $X(\boldsymbol{b}) := \mathrm{argmax}\{b_i : i \in [n]\}$. If there is one unique first-order bidder $|X(\boldsymbol{b})| = 1$, allocate the item to her $\mathrm{x}(\boldsymbol{b}) \equiv X(\boldsymbol{b})$. Otherwise $|X(\boldsymbol{b})| \geq 2$, allocate the item to one of those first-order bidders $\mathrm{x}(\boldsymbol{b}) \in X(\boldsymbol{b})$, via some (randomized) *tie-breaking* rule for this bid profile $\boldsymbol{b}$.
- **first-price payment:** The allocated bidder $\mathrm{x}(\boldsymbol{b})$ pays her own bid, while the non-allocated bidders $[n] \setminus \{\mathrm{x}(\boldsymbol{b})\}$ pay nothing. Formally, $\rho_i(\boldsymbol{b}) = b_i \cdot \mathbb{1}(i = \mathrm{x}(\boldsymbol{b}))$ for each $i \in [n]$.

Hence, different First Price Auctions $\mathcal{A} \in \mathbb{FPA}$ are identified by their allocation/tie-breaking rules $\mathrm{x}(\boldsymbol{b})$ and, without ambiguity, we can abuse the notation $\mathrm{x} \in \mathbb{FPA}$.

Regarding a joint value distribution $\boldsymbol{v} = (v_i)_{i \in [n]} \sim \boldsymbol{V} \in \mathbb{V}_{\mathsf{joint}}$, a (randomized) strategy profile $\boldsymbol{s} = \{s_i\}_{i \in [n]}$ maps the realized individual values v_i to the (random) individual bids $s_i(v_i)$. Over the randomness of other bidders' bids $\boldsymbol{s}_{-i}(\boldsymbol{v}_{-i})$ and the allocation rule $\mathrm{x} \in \mathbb{FPA}$, bidder $i \in [n]$ on having a value $v \geq 0$ and a bid $b \geq 0$ wins with probability $\mathrm{x}_i(b) := \mathbf{Pr}_{\boldsymbol{v},\, \boldsymbol{s},\, \mathrm{x}}[i = \mathrm{x}(b,\, \boldsymbol{s}_{-i}(\boldsymbol{v}_{-i})) \,|\, v_i = v]$ and gains an interim utility $u_i(v,\, b) := (v - b) \cdot \mathrm{x}_i(b)$. Such a strategy profile $\boldsymbol{s}$ forms a Bayesian Nash Equilibrium when it satisfies the following conditions.

DEFINITION 2.1. (BAYESIAN NASH EQUILIBRIA) *Given a joint value distribution $\boldsymbol{V} \in \mathbb{V}_{\mathsf{joint}}$, an allocation rule $\mathrm{x} \in \mathbb{FPA}$, and a precision $\delta > 0$:*

- *An (exact) Bayesian Nash Equilibrium $\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V},\, \mathrm{x})$ is a strategy profile $\boldsymbol{s} = \{s_i\}_{i \in [n]}$ that, for any bidder $i \in [n]$, any value of her $v \in \mathrm{supp}_i(\boldsymbol{V})$, and any deviation bid $b^* \geq 0$,*

$$\mathop{\mathbf{E}}_{\boldsymbol{v},\, \boldsymbol{s},\, \mathrm{x}}[\, u_i(v_i,\, \boldsymbol{s}(\boldsymbol{v})) \mid v_i = v\,] \;\geq\; \mathop{\mathbf{E}}_{\boldsymbol{v},\, \boldsymbol{s},\, \mathrm{x}}[\, u_i(v_i,\, b^*,\, \boldsymbol{s}_{-i}(\boldsymbol{v})) \mid v_i = v\,].$$

- *A δ-approximate Bayesian Nash Equilibrium $\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V},\, \mathrm{x})$ is a strategy profile $\boldsymbol{s} = \{s_i\}_{i \in [n]}$ that, for any bidder $i \in [n]$, any value of her $v \in \mathrm{supp}_i(\boldsymbol{V})$, and any deviation bid $b^* \geq 0$,*

$$\mathop{\mathbf{E}}_{\boldsymbol{v},\, \boldsymbol{s},\, \mathrm{x}}[\, u_i(v_i,\, \boldsymbol{s}(\boldsymbol{v})) \mid v_i = v\,] \;\geq\; \mathop{\mathbf{E}}_{\boldsymbol{v},\, \boldsymbol{s},\, \mathrm{x}}[\, u_i(v_i,\, b^*,\, \boldsymbol{s}_{-i}(\boldsymbol{v})) \mid v_i = v\,] - \delta.$$

2.1 Independent valuations. When the value distribution $\boldsymbol{V}$ degenerates into a product value distribution $\boldsymbol{V} = \{V_i\}_{i \in [n]} \in \mathbb{V}_{\mathsf{prod}}$, the equilibria thereof have several remarkable properties, which we give a brief review here.

First, the following result on the existence of exact equilibria can be concluded from [Leb96].

PROPOSITION 2.1. ([LEB96]) *Given a product value distribution $\boldsymbol{V} = \{V_i\}_{i \in [n]} \in \boldsymbol{V}_{\mathsf{prod}}$, there exists some tie-breaking rule $\mathrm{x} \in \mathbb{FPA}$ such that the resulting First Price Auction admits at least one exact equilibrium $\mathbb{BNE}(\boldsymbol{V},\, \mathrm{x}) \neq \emptyset$.*

Given an exact Bayesian Nash Equilibrium $\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V}, \mathrm{x})$, we will adopt the following notations.

- $\boldsymbol{B} = \{B_i\}_{i\in[n]}$ denotes the equilibrium bid distributions $\boldsymbol{s}(\boldsymbol{v}) = (s_i(v_i))_{i\in[n]} \sim \boldsymbol{B}$.
- $\mathcal{B}(b) = \prod_{i\in[n]} B_i(b)$ denotes the first-order bid distribution $\max(\boldsymbol{s}(\boldsymbol{v})) \sim \mathcal{B}$.
- $\mathcal{B}_{-i}(b) = \prod_{k\in[n]\setminus\{i\}} B_k(b)$ denotes the competing bid distribution of each bidder $i \in [n]$.
- $\gamma := \inf(\mathrm{supp}(\mathcal{B}))$ and $\lambda := \sup(\mathrm{supp}(\mathcal{B}))$ denote the "infimum"/"supremum" first-order bids, respectively. Without ambiguity, we call v, $b < \gamma$ the ***low*** values/bids, v, $b = \gamma$ the ***boundary*** values/bids, and v, $b > \gamma$ the ***normal*** values/bids. In other words: (i) low bids $b < \gamma$ give a zero winning probability and are less important; (ii) normal bids $b > \gamma$ are the most common bids and will behave nicely; and (iii) boundary bids $b = \gamma$ are tricky and will be dealt with separately.

The next proposition, due to [JL22, Lemma 2.7], shows that the equilibrium/competing/first-order bid distributions $B_i(b)$, $\mathcal{B}_{-i}(b)$, and $\mathcal{B}(b)$ have nice structures.

Proposition 2.2. ([JL22, Lemma 2.7]) *Each of the following holds:*

1. **monotonicity:** *The competing/first-order bid distributions $\{\mathcal{B}_{-i}\}_{i\in[n]}$ and $\mathcal{B}$ each have probability densities almost everywhere on $b \in (\gamma, \lambda]$, thus having strictly increasing CDF's on the closed interval $b \in [\gamma, \lambda]$.*
2. **continuity:** *The equilibrium/competing/first-order bid distributions $\{B_i\}_{i\in[n]}$, $\{\mathcal{B}_{-i}\}_{i\in[n]}$ and $\mathcal{B}$ each have no probability mass on $b \in (\gamma, \lambda]$, excluding the boundary $\gamma = \inf(\mathrm{supp}(\mathcal{B}))$, thus having continuous CDF's on the closed interval $b \in [\gamma, \lambda]$.*

Two more requisite notions for our later discussions are *bid-to-value mappings* and *monopolists* (Definitions 2.2 and 2.3). Particularly, we will leverage two structural results also from [JL22].

Definition 2.2. (Bid-to-value mappings) *The bid-to-value mappings $\boldsymbol{\varphi} = \{\varphi_i\}_{i\in[n]}$ are defined as $\varphi_i(b) := b + \mathcal{B}_{-i}(b)/\mathcal{B}'_{-i}(b) = b + \left(\sum_{k\in[n]\setminus\{i\}} B'_k(b)/B_k(b)\right)^{-1}$ for $b \in (\gamma, \lambda)$.*

Definition 2.3. (Monopolists) *A bidder $h \in [n]$ is called a* monopolist *when the probability of taking a normal value yet a boundary bid is nonzero $\mathbf{Pr}_{v_h,\, s_h}[(v_h > \gamma) \wedge (s_h(v_h) = \gamma)] > 0$.*

Proposition 2.3. ([JL22, Lemma 2.13]) *Each bid-to-value mapping $\varphi_i(b)$ for $i \in [n]$ is increasing on the open interval $b \in (\gamma, \lambda)$. Therefore, the domain can be extended to include the both endpoints $\varphi_i(\gamma) := \lim_{b\searrow\gamma} \varphi_i(b)$ and $\varphi_i(\lambda) := \lim_{b\nearrow\lambda} \varphi_i(b)$.*

Proposition 2.4. ([JL22, Lemma 2.16]) *There exists at most one monopolist $h \in [n]$. If existential:*
(I) *A boundary first-order bid $\{\max(\boldsymbol{b}) = \gamma\}$ occurs with a nonzero probability $\mathcal{B}(\gamma) > 0$.*
(II) *Conditioned on the tiebreak $\{b_h = \max(\boldsymbol{b}) = \gamma\}$, the monopolist wins $\mathrm{x}(\boldsymbol{b}) = h$ almost surely.*

3 Tie-breaking Rules and Approximate Equilibria.

In this section, we discuss the existence of exact/approximate equilibria for a product value distribution $\boldsymbol{V} = \{V_i\}_{i\in[n]} \in \mathbb{V}_{\mathsf{prod}}$. Following Proposition 2.1, the only possibility for nonexistence of exact equilibria is that the underlying tie-breaking rule $\mathrm{x} \in \mathbb{FPA}$ may be incompatible with this value distribution. To better understand this, let us give a simple example. Consider uniform tie-breaking rule and two agents: Alice has a deterministic value 1, and Bob has a deterministic value 0. In this example, Bob always bids 0. If Alice also bids $b = 0$, she gets the item with probability 50% under the uniform tie-breaking rule. This is not an equilibrium, since Alice can win the item with probability 100% by slightly increasing her bid. On the other hand, any nonzero bid $b > 0$ also is sub-optimal, since a lower bid, say $b' = b/2 > 0$, still ensures Alice to win. We think the tie-breaking issue is incurred by mathematical formulation, rather than the essence of equilibria. In real business, there is usually a minimum unit for bidding, say 0.01\$. Then "Alice bids 0.01 and Bob bids 0" is a Nash equilibrium under the uniform tie-breaking rule, since there is no other bid between 0 and 0.01. To treat this formally, we introduce the notion of approximate equilibrium.

We will start with a *compatible* tie-breaking and an *exact* equilibrium thereof $\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V}, \mathrm{x})$. Then for any given $\delta > 0$, we slightly modify this equilibrium into a new strategy profile $\boldsymbol{s}^*$ that is insensitive to different tie-breaking rules. I.e., this strategy profile $\boldsymbol{s}^*$ is a *universal* δ-approximate equilibrium for First Price Auction, regardless of the tie-breaking rules.

To make the modification workable, we crucially leverage several structural results from [JL22] about Bayesian Nash Equilibrium. In particular, we will use the fact that nontrivial tie-breaks can occur only when the first-order bid is at the boundary $\gamma = \inf(\mathrm{supp}(\mathcal{B}))$.

Before giving the formal statement of our result, we recall the concept of *earth mover's distance* [Vil09, Chapter 6], which will be used to measure the distance between two strategies.

Definition 3.1. (Earth Mover's Distance) *Given two single-dimensional distributions D and $\widetilde{D}$, denote by $D^{-1}(q)$ and $\widetilde{D}^{-1}(q)$ for $q \in [0, 1]$ the quantile functions, then:*

- *The ℓ_p-norm earth mover's distance, for $p \geq 1$, is defined as*

$$\mathsf{EMD}_p(D, \widetilde{D}) \;=\; \Big(\int_0^1 \big|D^{-1}(q) - \widetilde{D}^{-1}(q)\big|^p \cdot \mathrm{d}q\Big)^{1/p}.$$

- *The ℓ_∞-norm earth mover's distance is defined as*

$$\mathsf{EMD}_\infty(D, \widetilde{D}) \;=\; \sup\Big\{\big|D^{-1}(q) - \widetilde{D}^{-1}(q)\big| : q \in [0, 1]\Big\}.$$

It follows that $\mathsf{EMD}_{p_1}(D, \widetilde{D}) \leq \mathsf{EMD}_{p_2}(D, \widetilde{D}) \leq \mathsf{EMD}_\infty(D, \widetilde{D})$ for any $p_2 \geq p_1 \geq 1$.

Below, Theorem 3.1 summarizes our result on the existence of universal δ-approximate equilibria. The proof relies on [JL22, Lemma 2.5].

Proposition 3.1. ([JL22, Lemma 2.5]) *At an exact Bayesian Nash Equilibrium $\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V}, \mathrm{x})$, for each bidder $i \in [n]$, the following hold almost surely:*

1. *A low/boundary value $v \in \mathrm{supp}_{\leq\gamma}(V_i)$ induces a low/boundary equilibrium bid $s_i(v) \leq \gamma$.*
2. *A normal value $v \in \mathrm{supp}_{>\gamma}(V_i)$ induces a boundary/normal equilibrium bid $\gamma \leq s_i(v) < v$.*

Theorem 3.1. (Bayesian Nash Equilibria) *Given a product value distribution $\boldsymbol{V} = \{V_i\}_{i\in[n]} \in \mathbb{V}_{\mathsf{prod}}$, a tie-breaking rule $\mathrm{x} \in \mathbb{FPA}$, and an exact Bayesian Nash Equilibrium thereof $\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V}, \mathrm{x}) \neq \emptyset$. For any precision $\delta > 0$, there exists another strategy profile $\boldsymbol{s}^* = \{s_i^*\}_{i\in[n]}$ such that:*

1. **closeness:** $\mathsf{EMD}_\infty(s_i(v), s_i^*(v)) \leq \delta$ *for any value $v \in \mathrm{supp}(V_i)$ and each bidder $i \in [n]$.*[2]
2. **efficiency invariant:** *For an arbitrary tie-breaking rule $\mathrm{x}^* \in \mathbb{FPA}$ (possibly the same as x), the expected optimal/auction Social Welfares keep the same $\mathsf{OPT}(\boldsymbol{V}, \mathrm{x}^*, \boldsymbol{s}^*) = \mathsf{OPT}(\boldsymbol{V}, \mathrm{x}, \boldsymbol{s})$ and $\mathsf{FPA}(\boldsymbol{V}, \mathrm{x}^*, \boldsymbol{s}^*) = \mathsf{FPA}(\boldsymbol{V}, \mathrm{x}, \boldsymbol{s})$.*
3. **universality:** *For an arbitrary tie-breaking rule $\mathrm{x}^* \in \mathbb{FPA}$ (possibly the same as x), it forms a δ-approximate equilibrium $\boldsymbol{s}^* \in \mathbb{BNE}(\boldsymbol{V}, \mathrm{x}^*, \delta)$. Formally, it forms a universal δ-approximate equilibrium $\boldsymbol{s}^* \in \big(\bigcap_{\mathrm{x}^*\in\mathbb{FPA}} \mathbb{BNE}(\boldsymbol{V}, \mathrm{x}^*, \delta)\big)$.*

Proof. There are 9 kinds of tuples $(v_i, s_i(v_i))$, i.e., low/boundary/normal values v_i and bids $s_i(v_i)$. Based on case analysis, we construct the new strategy profile $\boldsymbol{s}^* = \{s_i^*\}_{i\in[n]}$ in a *coupling* way.

[2] Recall that each strategy s_i or s_i^* is a family of bid distributions indexed by the value $v \in \mathrm{supp}(V_i)$.

value \ bid	low $s_i(v_i) < \gamma$	BDY $s_i(v_i) = \gamma$	normal $s_i(v_i) > \gamma$
low $v_i < \gamma$	$s_i^*(v_i) = s_i(v_i) - \delta$	$s_i^*(v_i) = s_i(v_i) - \delta$	impossible (Prop. 3.1)
BDY $v_i = \gamma$	$s_i^*(v_i) = s_i(v_i) - \delta$	$s_i^*(v_i) = s_i(v_i) - \delta/2$	impossible (Prop. 3.1)
normal $v_i > \gamma$	impossible (Prop. 3.1)	$s_i^*(v_i) = s_i(v_i)$	$s_i^*(v_i) = s_i(v_i)$

This coupling intrinsically ensures **Part 1** that $\mathsf{EMD}_\infty(s_i(v),\, s_i^*(v)) \le \delta$ for any value $v \in \mathrm{supp}(V_i)$ and each bidder $i \in [n]$. It remains to show **Part 2** and **Part 3**.

Part 2. The expected optimal $\mathsf{Social\ Welfare}$, which relies just on the value distribution $\boldsymbol{V}$, must be invariant $\mathsf{OPT}(\boldsymbol{V},\, \mathrm{x}^*,\, \boldsymbol{s}^*) = \mathsf{OPT}(\boldsymbol{V},\, \mathrm{x},\, \boldsymbol{s})$. To reason about the expected auction $\mathsf{Social\ Welfare}$, recall Proposition 2.4 that the allocated bidder $\mathrm{x} = \mathrm{x}(\boldsymbol{s}(\boldsymbol{v}))$ has three possibilities:

- **Case (I).** The allocated bidder x has a normal bid $s_{\mathrm{x}}(v_{\mathrm{x}}) > \gamma$ and a normal value $v_{\mathrm{x}} > \gamma$.

  Regarding the coupling between both strategy profiles $\boldsymbol{s}^* = \{s_i^*\}_{i\in[n]}$ and $\boldsymbol{s} = \{s_i\}_{i\in[n]}$, normal bidders $\{i \in [n] \,|\, s_i(v_i) > \gamma\}$ preserve their bids $s_i^*(v_i) = s_i(v_i)$, while low/boundary bidders $\{j \in [n] \,|\, s_j(v_j) \le \gamma\}$ never increase their bids $s_j^*(v_j) \le s_j(v_j)$.

  In the coupled scenario $(\mathrm{x}^*,\, \boldsymbol{s}^*(\boldsymbol{v}))$, bidder x still has the first-order bid $s_{\mathrm{x}}^*(v_{\mathrm{x}}) = \max(\boldsymbol{s}^*(\boldsymbol{v}))$ and, (Item 2 of Proposition 2.2) almost surely, is the only first-order bidder $\mathrm{argmax}(\boldsymbol{s}^*(\boldsymbol{v})) = \{\mathrm{x}\}$ and thus keeps winning $\mathrm{x}^*(\boldsymbol{s}^*(\boldsymbol{v})) = \mathrm{x}$.

- **Case (II).** The allocated bidder x has a boundary bid $s_{\mathrm{x}}(v_{\mathrm{x}}) = \gamma$ and a normal value $v_{\mathrm{x}} > \gamma$.

  Bidder x is the *unique* monopolist (Proposition 2.4). Further, other bidders $i \in [n] \setminus \{\mathrm{x}\}$ have low/boundary bids $s_i(v_i) \le \gamma$ and low/boundary values $v_i \le \gamma$ (cf. the $\boldsymbol{s}^*$-construction table).

  In the coupled scenario $(\mathrm{x}^*,\, \boldsymbol{s}^*(\boldsymbol{v}))$, bidder x preserves her bid $s_{\mathrm{x}}^*(v_{\mathrm{x}}) = s_{\mathrm{x}}(v_{\mathrm{x}}) = \gamma$, while other bidders $i \in [n] \setminus \{\mathrm{x}\}$ decrease their bids $s_i^*(v_i) \le s_i(v_i) - \delta/2 \le \gamma - \delta/2$. Thus, bidder x is the only first order bidder $\mathrm{argmax}(\boldsymbol{s}^*(\boldsymbol{v})) = \{\mathrm{x}\}$ and thus keeps winning $\mathrm{x}^*(\boldsymbol{s}^*(\boldsymbol{v})) = \mathrm{x}$.

- **Case (III).** The allocated bidder x has a boundary bid/value $s_{\mathrm{x}}(v_{\mathrm{x}}) = v_{\mathrm{x}} = \gamma$.

  The original scenario has *no* monopolist (Proposition 2.4), almost surely. All bidders $i \in [n]$ have low/boundary bids $s_i(v_i) \le \gamma$ and low/boundary values $v_i \le \gamma$ (cf. the $\boldsymbol{s}^*$-construction table). Further, bidders $N_{\mathrm{BDY}} = \{j \in [n] \,|\, s_j(v_j) = v_j = \gamma\}$, including the allocated bidder x, have boundary bids/values.

  In the coupled scenario $(\mathrm{x}^*,\, \boldsymbol{s}^*(\boldsymbol{v}))$, bidders $j \in N_{\mathrm{BDY}}$ have the bid $s_j^*(v_j) = s_j(v_j) - \delta/2 = \gamma - \delta/2$, while the other bidders $i \in [n] \setminus N_{\mathrm{BDY}}$ have lower bids $s_i^*(v_i) = s_i(v_i) - \delta < \gamma - \delta$. Thus, bidders $j \in N_{\mathrm{BDY}}$ are exactly the coupled first-order bidders $\mathrm{argmax}(\boldsymbol{s}^*(\boldsymbol{v})) = N_{\mathrm{BDY}}$. I.e., regardless of the tie-breaking rule, any possible allocation $\mathrm{x}^* = \mathrm{x}^*(\boldsymbol{s}^*(\boldsymbol{v})) \in N_{\mathrm{BDY}}$ always realizes the boundary $\mathsf{Social\ Welfare}$ $v_{\mathrm{x}^*} = \gamma$, the same as the original scenario $v_{\mathrm{x}} = \gamma$.

So, the coupling between $\boldsymbol{s}^* = \{s_i^*\}_{i\in[n]}$ and $\boldsymbol{s} = \{s_i\}_{i\in[n]}$ preserves the same auction $\mathsf{Social\ Welfare}$, almost surely. In expectation, we have $\mathsf{FPA}(\boldsymbol{V},\, \mathrm{x}^*,\, \boldsymbol{s}^*) = \mathsf{FPA}(\boldsymbol{V},\, \mathrm{x},\, \boldsymbol{s})$.

Part 3. Under our coupling: The normal bids $s_i(v_i) > \gamma$ keep the same $s_i^*(v_i) = s_i(v_i)$. The low bids $s_i(v_i) < \gamma$ are shifted by a $-\delta$ distance, namely $s_i^*(v_i) = s_i(v_i) - \delta$. Moreover, the boundary bids $s_i(v_i) = \gamma$ are "split" into (i) $s_i^*(v_i) = \gamma - \delta$ for low values $v_i < \gamma$, (ii) $s_i^*(v_i) = \gamma - \delta/2$ for boundary values $v_i = \gamma$, or (iii) $s_i^*(v_i) = \gamma$ for high values $v_i > \gamma$, which occurs only for the unique monopolist (if existential).

The coupled infimum first-order bid $\gamma^* = \inf(\mathrm{supp}(\mathcal{B}^*))$ is bounded between $\gamma^* \in [\gamma - \delta/2,\, \gamma]$, since in the original scenario, conditioned on the boundary first-order bid $\{\max(\boldsymbol{s}(\boldsymbol{v})) = \gamma\}$, there always exists at least one boundary/normal valuer $\{\max(\boldsymbol{v}) \ge \gamma\}$ (Proposition 2.4). We would verify the δ-approximate equilibrium conditions, through on case analysis about the original scenario:

- A low/boundary value/bid $v_i \le \gamma$ and $s_i(v_i) \le \gamma$, with at least one strict inequality.

By construction, the coupled bid $s_i^*(v_i) \le \gamma - \delta$ is strictly below the coupled infimum first-order bid $\gamma^* \in [\gamma - \delta/2,\, \gamma]$, yielding a zero interim allocation/utility $= 0$. In contrast, because this bidder i has a low/boundary value $v_i \le \gamma$, any deviation bid $b^* \ge 0$ yields an interim utility at most $\le \min(v_i - \gamma^*,\, 0) \le \delta/2 < \delta$.

- A boundary value/bid $v_i = s_i(v_i) = \gamma$.

  By construction, the coupled bid $s_i^*(v_i) = \gamma - \delta/2 < v_i = \gamma$ yields a nonnegative interim utility ≥ 0. In contrast, because this bidder i has a boundary value $v_i = \gamma$, any deviation bid $b^* \ge 0$ yields an interim utility at most $\le v_i - \gamma^* = \delta/2 < \delta$.

- A normal value $v_i > \gamma$ and a boundary/normal bid $s_i(v_i) \ge \gamma$.

  Following Item 2 of Proposition 2.2, the coupled bid $s_i^*(v_i) = s_i(v_i)$ yields the same nonnegative interim allocation/utility ≥ 0 as in the original scenario $(\mathrm{x},\, \boldsymbol{s}(\boldsymbol{v}))$. In comparison: (i) Any deviation bid $b^* \ge \gamma$ yields a smaller or equal interim utility, as a consequence of the exact equilibrium $\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V},\, \mathrm{x})$. (ii) Any deviation bid $b^* < \gamma^*$ yields a zero interim utility $= 0$. (iii) Any deviation bid $b^* \in [\gamma^*,\, \gamma)$ yields at most "the bid-γ interim utility $\le$ the current interim utility" plus "a term of $\gamma - b^* \le \gamma - \gamma^* = \delta/2 < \delta$".

Hence, the coupled strategy profile forms a δ-approximate equilibrium $\boldsymbol{s}^* \in \mathbb{BNE}(\boldsymbol{V},\, \mathrm{x}^*,\, \delta)$.

A minor issue is the above modification may incur negative bids if the original infimum first-order bid is too small $\gamma < \delta$. Instead, we can first slightly shift the original strategies $\boldsymbol{s} = \{s_i\}_{i \in [n]}$ by a $+\delta$ distance and then reapply the above modification. As a consequence, everything keeps the same, except that the bidders' utilities each drop by a δ amount. This finishes the proof. ∎

A revelation of Theorem 3.1 is that we can focus on *exact* equilibria $\boldsymbol{s}$ in studying the PoA/PoS problems, as if we can control the tie-breaking rule x and choose a *compatible* one $\mathbb{BNE}(\boldsymbol{V},\, \mathrm{x}) \ne \emptyset$. When an *incompatible* tie-breaking rule x^* are really considered, up to any precision $\delta > 0$, we can still obtain a δ-approximate equilibrium $\boldsymbol{s}^* \in \mathbb{BNE}(\boldsymbol{V},\, \mathrm{x}^*,\, \delta)$ by modifying any "compatible" exact equilibrium $\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V},\, \mathrm{x})$. We will adopt this convention in Sections 4 and 5, since it simplifies the notation and (essentially) incurs no loss of generality.

4 Bayesian Nash Equilibria for Independent Valuations.

In this section, we prove the following tight PoS result for the canonical setting, namely Bayesian Nash Equilibria for independent valuations.

Theorem 4.1. (Tight PoS) *Regarding Bayesian Nash Equilibria for independent valuations, the Price of Stability is* $1 - 1/e^2 \approx 0.8647$.

The lower-bound part of Theorem 4.1 immediately follows from the tight PoA result from [JL22]. To get the upper-bound part, we only need to construct an instance whose PoS is exactly $1 - 1/e^2$. Technically, we provide a sequence of instances whose PoS asymptotically approaches $1 - 1/e^2$. The following instances are a slight modification from the tight PoA instances due to [JL22, Example 4] such that each modified instance has one unique equilibrium.

Example. Given an arbitrarily small constant $\varepsilon \in (0,\, 1/8)$, consider the $(n+1)$-bidder instance $\{H\} \cup \{L_i\}_{i \in [n]}$ for $n = \lceil 1/\varepsilon \rceil \ge 8$ in terms of value distributions $\boldsymbol{V} = V_H \otimes \{V_L\}^{\otimes n}$.

- Bidder H has a Bernoulli random value $v_H \sim V_H$ that $\mathbf{Pr}[v_H = 0] = \varepsilon$ and $\mathbf{Pr}[v_H = 1] = 1 - \varepsilon$.
- Bidders $\{L_i\}_{i \in [n]}$ have i.i.d. values $(v_{L,\,i})_{i \in [n]} \sim \{V_L\}^{\otimes n}$ whose common value distribution V_L is given by the parametric equation $V_L(1 - \frac{n-(t-1)}{nt-(t-1)} \cdot t^2 \cdot e^{2-2t}) = \sqrt[n]{4/t^2 \cdot e^{2t-4}}$ for $t \in [1,\, 2]$. This value distribution V_L is supported on $\mathrm{supp}(V_L) = [0,\, 1 - \frac{2n-2}{2n-1} \cdot 2/e^2]$ and has a probability mass $V_L(0) = \sqrt[n]{4/e^2}$ at the zero value.

The considered First Price Auction $\mathrm{x} \in \mathbb{FPA}$, under the all-zero bid profile $\boldsymbol{b} = b_H \otimes (b_{L,\,i})_{i \in [n]} = \mathbf{0}$, favors bidder H, but otherwise is arbitrary. ∎

4.1 The focal equilibrium. In this part, we verify that the worst-case equilibrium for the tight PoA instance [JL22, Example 4], after a slight adjustment, is still an equilibrium for our modified instance. (The proof is similar to the equilibrium condition analysis in [JL22, Section 6].)

For clarity, this modified equilibrium will be called the ***focal*** strategy profile or, after verifying the equilibrium condition, the ***focal*** equilibrium. Let $\lambda^* := 1 - 4/e^2 \approx 0.4587$. The focal strategy profile $\boldsymbol{s}^* = \{s_H^*\} \otimes \{s_L^*\}^{\otimes n}$ is given as follows; see Figure 1 for a visual aid.

- Bidder H has a (mixed) strategy $s_H^*(0) \equiv 0$ for a zero value $\{v_H = 0\}$ and $s_H^*(1) \sim S_H^*$ for a nonzero value $\{v_H = 1\}$. Here the bid distribution S_H^* is given by the implicit equation $b = 1 - 4 \cdot (\varepsilon + (1-\varepsilon) \cdot S_H^*) \cdot e^{2-4\sqrt{\varepsilon + (1-\varepsilon)\cdot S_H^*}}$ for $S_H^* \in [1 - \frac{3/4}{1-\varepsilon},\, 1]$.

  The random bid $s_H^*(v_H)$ is supported on $\{0\} \cup \mathrm{supp}(S_H^*) = [0,\, 1 - 4/e^2] = [0,\, \lambda^*]$.

- Bidders $\{L_i\}_{i\in[n]}$ have (deterministic) identical strategies $\{s_L^*\}^{\otimes n}$ that are given by the parametric equation $s_L^*(1 - \frac{n-(t-1)}{nt-(t-1)} \cdot t^2 \cdot e^{2-2t}) = 1 - t^2 \cdot e^{2-2t}$ for $t \in [1,\, 2]$.

  The random bids $s_L^*(v_{L,\,i})$ are supported on $\{1 - t^2 \cdot e^{2-2t} \,|\, t \in [1,\, 2]\} = [0,\, 1 - 4/e^2] = [0,\, \lambda^*]$.

Lemma 4.1 checks the equilibrium condition for the focal strategy profile $\boldsymbol{s}^*$.

Lemma 4.1. (Equilibrium) *The following hold for the focal strategy profile $\boldsymbol{s}^* = \{s_H^*\} \otimes \{s_L^*\}^{\otimes n}$:*

1. *Bidder H has the bid distribution B_H^* given by the implicit equation $b = 1 - 4B_H^* \cdot e^{2-4\sqrt{B_H^*}}$ for $B_H^* \in [1/4,\, 1]$, and a constant bid-to-value mapping $\varphi_H(b) = 1$ for $b \in [0,\, \lambda^*]$.*

2. *Bidders $\{L_i\}_{i\in[n]}$ have identical bid distributions $\{B_L^*\}^{\otimes n}$ given by $B_L^*(b) = \sqrt[n]{(1-\lambda^*)/(1-b)}$ for $b \in [0,\, \lambda^*]$, and identical bid-to-value mappings $\{\varphi_L^*\}^{\otimes n}$ given by the parametric equation $\varphi_L^*(1 - t^2 \cdot e^{2-2t}) = 1 - \frac{n-(t-1)}{nt-(t-1)} \cdot t^2 \cdot e^{2-2t}$ for $t \in [1,\, 2]$.*

3. *The focal strategy profile $\boldsymbol{s}^* = s_H^* \otimes \{s_L^*\}^{\otimes n}$ forms a Bayesian Nash Equilibrium $\boldsymbol{s}^* \in \mathbb{BNE}(\boldsymbol{V})$.*

Proof. We first reason about **Item 1** and **Item 2**.

For bidder H, the strategy s_H^* converts (i) all densities at the zero value $\mathbf{Pr}[v_H = 0] = \varepsilon$ to densities at the zero bid $s_H^*(0) \equiv 0$ and (ii) all densities at the nonzero value $\mathbf{Pr}[v_H = 1] = 1 - \varepsilon$ to densities that follow the bid distribution S_H^*, which is supported on $\mathrm{supp}(S_H^*) = [0,\, \lambda^*]$. Overall, the bid distribution $s_H^*(v_H) \sim B_H^*$ can be written as $B_H^*(b) = \varepsilon + (1-\varepsilon) \cdot S_H^*(b)$, over the bid support $b \in [0,\, \lambda^*]$. Plugging this formula into the defining implicit equation for S_H^*, we can conclude with $b = 1 - 4B_H^* \cdot e^{2-4\sqrt{B_H^*}}$ for $B_H^* \in [1/4,\, 1]$, as desired.

For bidders $\{L_i\}_{i\in[n]}$ and their identical strategies $\{s_L^*\}^{\otimes n}$, the value formula $1 - \frac{n-(t-1)}{nt-(t-1)} \cdot t^2 \cdot e^{2-2t}$ and the bid formula $1 - t^2 \cdot e^{2-2t}$ both are increasing in $t \in [1,\, 2]$. Thus, the identical bid distributions $\{B_L^*\}^{\otimes n}$ can be written as $\{(b,\, B_L^*) = (s_L^*(v),\, V_L(v)) \,|\, v \in \mathrm{supp}(V_L) = [0, 1 - \frac{2n-2}{2n-1} \cdot 2/e^2]\}$. Plugging the defining parametric equations for s_L^* and V_L into this formula, those bid distributions $\{B_L^*\}^{\otimes n}$ can be formulated as $\{(b,\, B_L^*) = (1 - t^2 \cdot e^{2-2t},\ \sqrt[n]{4/t^2 \cdot e^{2t-4}}) \,|\, t \in [1,\, 2]\}$. After rearranging, we can conclude with $B_L^*(b) = \sqrt[n]{(1-\lambda^*)/(1-b)}$ for $b \in [0,\, \lambda^*]$, as desired.

Bidder H competes with bidders $\{L_i\}_{i\in[n]}$, thus having the competing bid distribution $\mathcal{B}_{-H}^*(b) = \big(B_L^*(b)\big)^n = (1-\lambda^*)/(1-b)$ and the *constant* bid-to-value mapping $\varphi_H^*(b) = b + \mathcal{B}_{-H}^*(b)/{\mathcal{B}_{-H}^*}'(b) = 1$ for $b \in [0,\, \lambda^*]$, as desired.

Let $t := 2\sqrt{B_H^*} \in [1,\, 2]$. Then we have $b = 1 - 4B_H^* \cdot e^{2-4\sqrt{B_H^*}} = 1 - t^2 \cdot e^{2-2t}$ and the derivative $\frac{\mathrm{d}b}{\mathrm{d}t} = (2t^2 - 2t) \cdot e^{2-2t}$. Each bidder L_i for $i \in [n]$ competes with bidders $\{H\} \cup \{L_j\}_{j\in[n]\setminus\{i\}}$, hence the competing bid distribution $\mathcal{B}_{-L}^*(b) = B_H^*(b) \cdot \big(B_L^*(b)\big)^{n-1}$. In terms of the parameter $t \in [1,\, 2]$, we can substitute $B_H^* = t^2/4$, $\lambda^* = 1 - 4/e^2$, and $b = 1 - t^2 \cdot e^{2-2t}$, rewriting

$$\mathcal{B}_{-L}^* \;=\; B_H^* \cdot \big(B_L^*\big)^{n-1} \;=\; B_H^* \cdot \Big(\frac{1-\lambda^*}{1-b}\Big)^{\frac{n-1}{n}} \;=\; t^2/4 \cdot \Big(\frac{4/t^2}{e^{4-2t}}\Big)^{\frac{n-1}{n}}.$$

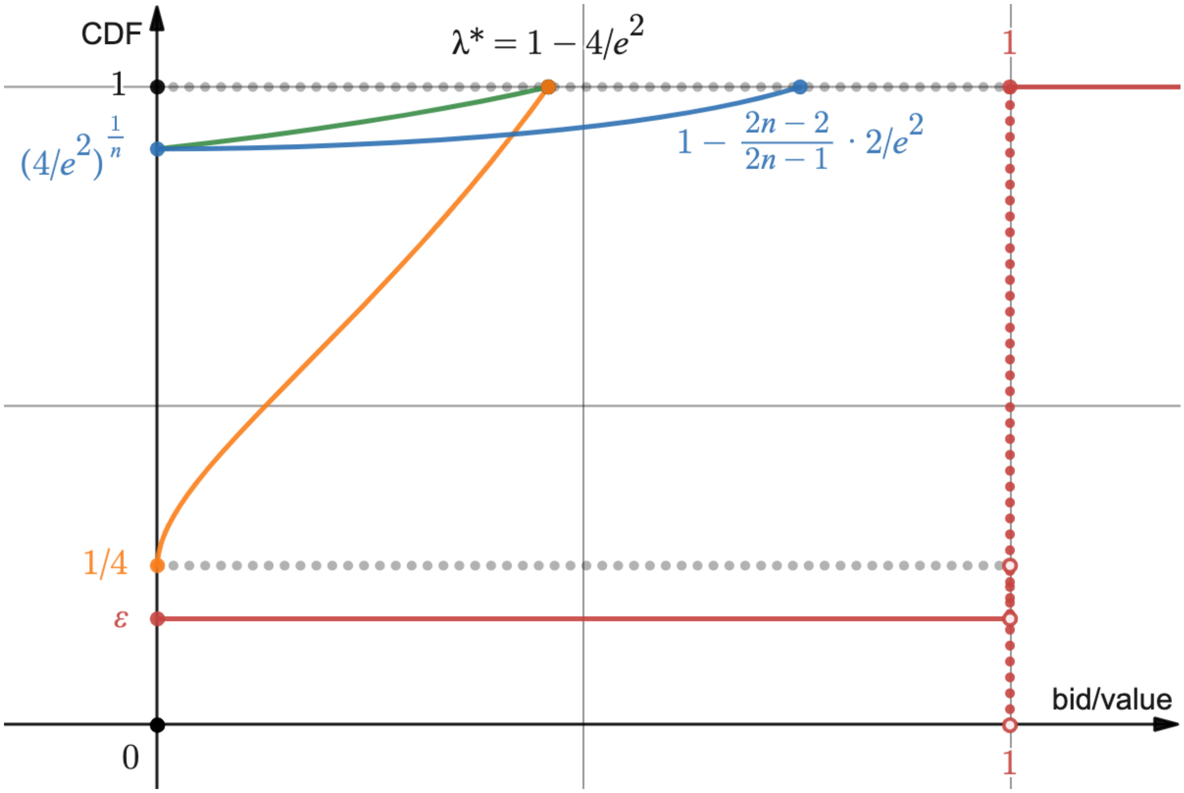

Figure 1: Demonstration for the $(n+1)$-bidder instance $\boldsymbol{V} = V_H \otimes \{V_L\}^{\otimes n}$ in Section 4.
(red) Bidder H has a Bernoulli value distribution $V_H(v) = \varepsilon$ for $v \in [0,\ 1)$ and $V_H(v) = 1$ for $v \geq 1$.
(blue) Bidders $\{L_i\}_{i \in [n]}$ have identical value distributions $\{V_L\}^{\otimes n}$ given by the parametric equation $V_L(1 - \frac{n-(t-1)}{nt-(t-1)} \cdot t^2 \cdot e^{2-2t}) = \sqrt[n]{4/t^2 \cdot e^{2t-4}}$ for $t \in [1,\ 2]$.
Under the focal strategy profile $\boldsymbol{s}^* = \{s_H^*\} \otimes \{s_L^*\}^{\otimes n}$: Over the support $b \in [0,\ \lambda^*]$, the resulting bid distributions $\boldsymbol{B}^* = B_H^* \otimes \{B_L^*\}^{\otimes n}$ are given by (orange) the implicit equation $b = 1 - 4B_H^* \cdot e^{2-4\sqrt{B_H^*}}$ for $B_H^* \in [1/4,\ 1]$ and (green) the parametric equation $B_L^*(1 - t^2 \cdot e^{2-2t}) = \sqrt[n]{4/t^2 \cdot e^{2t-4}}$ for $t \in [1,\ 2]$.

Then in terms of $t \in [1,\ 2]$, the bid-to-value mapping $\varphi_L^*(b) = b + \frac{\mathcal{B}_{-L}^*}{\mathrm{d}\mathcal{B}_{-L}^*/\mathrm{d}b}$ is given by

$$\begin{aligned}\varphi_L^* &= b + \mathcal{B}_{-L}^* \cdot \frac{\mathrm{d}b/\mathrm{d}t}{\mathrm{d}\mathcal{B}_{-L}^*/\mathrm{d}t} \\ &= (1 - t^2 \cdot e^{2-2t}) + nt \cdot \frac{(t^2 - t) \cdot e^{2-2t}}{nt - (t-1)} \\ &= 1 - \frac{n - (t-1)}{nt - (t-1)} \cdot t^2 \cdot e^{2-2t}.\end{aligned}$$

Hence, we obtain the parametric equation $\varphi_L^*(1 - t^2 \cdot e^{2-2t}) = 1 - \frac{n-(t-1)}{nt-(t-1)} \cdot t^2 \cdot e^{2-2t}$ for $t \in [1,\ 2]$, as desired. Notably, this bid-to-value mapping is the inverse function of $\{L_i\}_{i\in[n]}$'s focal strategies $s_L^*(1 - \frac{n-(t-1)}{nt-(t-1)} \cdot t^2 \cdot e^{2-2t}) = 1 - t^2 \cdot e^{2-2t}$ for $t \in [1,\ 2]$.[3] **Item 1** and **Item 2** follow then.

Item 3. Bidder H has the interim utility formula $u_H^*(v_H,\ b) = (v_H - b) \cdot \mathcal{B}_{-H}^*(b) = (v_H - b) \cdot \frac{1-\lambda^*}{1-b}$ for $b \in [0,\ \lambda^*]$. Under a zero value $\{v_H = 0\}$, clearly the focal bid $s_H^*(0) \equiv 0$ must be utility-optimal. Under a nonzero value

[3] It is easy to check that both formulas $1 - \frac{n-(t-1)}{nt-(t-1)} \cdot t^2 \cdot e^{2-2t}$ and $1 - t^2 \cdot e^{2-2t}$ are strictly increasing in $t \in [1,\ 2]$.

$\{v_H = 1\}$, all bids $b \in [0, \lambda^*]$ yield the same interim utility $u_H^*(1, b) = 1 - \lambda^*$, so the focal bid $s_H^*(1) \sim S_H^*$ that $\mathrm{supp}(S_H^*) = [0, \lambda^*]$ also is utility-optimal.

Moreover, bidders $\{L_i\}_{i \in [n]}$ have the same interim utility formula $u_L^*(v_L, b) = (v_L - b) \cdot \mathcal{B}_{-L}^*(b)$ for $b \in [0, \lambda^*]$. For any given value $v_L \in [0, 1 - \frac{2n-2}{2n-1} \cdot 2/e^2]$, a bid $b \in [0, \lambda^*]$ is utility-optimal when it satisfies that $0 = \frac{\partial}{\partial b} u_L^*(v_L, b) = -\mathcal{B}_{-L}^*(b) + (v_L - b) \cdot \mathcal{B}_{-L}^{*\prime}(b)$, or equivalently, that $v_L = \varphi_L^*(b)$. The focal bid $s_L^*(v_L)$ *is* utility optimal, namely $v_L = \varphi_L^*(s_L^*(v_L))$, because the bid-to-value mapping φ_L^* is the inverse function of the focal strategy s_L^*.

Thus, all bidders $\{H\} \cup \{L_i\}_{i \in [n]}$ meet the equilibrium conditions. **Item 3** follows then. ∎

Below, Lemma 4.3 measures the expected optimal/auction Social Welfares from our $(n+1)$-bidder instance $\boldsymbol{V} = V_H \otimes \{V_L\}^{\otimes n}$ at the focal equilibrium $\boldsymbol{s}^* = \{s_H^*\} \otimes \{s_L^*\}^{\otimes n}$. The proof relies on the auction Social Welfare formula from [JL22, Lemma 2.20].

LEMMA 4.2. (AUCTION SOCIAL WELFARE [JL22]) *The expected auction Social Welfare $\mathsf{FPA}(\boldsymbol{V}, \boldsymbol{s})$ at a Bayesian Nash Equilibrium $\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V})$, on having a monopolist H, can be formulated as follows:*

$$\mathsf{FPA}(\boldsymbol{V}, \boldsymbol{s}) \;=\; \mathop{\mathbf{E}}_{v_H,\, s_H}[v_H \mid s_H(v_H) = \gamma] \cdot \mathcal{B}(\gamma) + \sum_{i \in [n]} \bigg(\int_\gamma^\lambda \varphi_i(b) \cdot \frac{B_i'(b)}{B_i(b)} \cdot \mathcal{B}(b) \cdot \mathrm{d}b \bigg).$$

LEMMA 4.3. (EFFICIENCY) *The following hold for the focal equilibrium $\boldsymbol{s}^* = \{s_H^*\} \otimes \{s_L^*\}^{\otimes n}$:*

1. *The expected optimal Social Welfare $\mathsf{OPT}(\boldsymbol{V}, \boldsymbol{s}^*) \geq 1 - \varepsilon$.*
2. *The expected auction Social Welfare $\mathsf{FPA}(\boldsymbol{V}, \boldsymbol{s}^*) \leq 1 - (1 - \varepsilon) \cdot e^{-2}$.*

Proof. Let us prove **Items 1 and 2** one by one.

Item 1. The realized optimal Social Welfare is at least bidder H's realized value $v_H \sim V_H$, namely a Bernoulli random value $\mathbf{Pr}[v_H = 0] = \varepsilon$ and $\mathbf{Pr}[v_H = 1] = 1 - \varepsilon$ (Section 4). In expectation, we have $\mathsf{OPT}(\boldsymbol{V}, \boldsymbol{s}^*) \geq \mathbf{E}[v_H] = 1 - \varepsilon$. **Item 1** follows then.

Item 2. Following Lemma 4.2, with the focal first-order bid CDF $\mathcal{B}^*(b) = B_H^*(b) \cdot \big(B_L^*(b)\big)^n$, the expected auction Social Welfare $\mathsf{FPA}(\boldsymbol{V}, \boldsymbol{s}^*)$ from our $(n+1)$-bidder instance is given by

$$\begin{aligned}
&\mathsf{FPA}(\boldsymbol{V}, \boldsymbol{s}^*) \\
=\;& \mathop{\mathbf{E}}_{v_H,\, s_H^*}[v_H \mid s_H^*(v_H) = 0] \cdot \mathcal{B}^*(0) + \int_0^{\lambda^*} \Big(\varphi_H^*(b) \cdot \frac{B_H^{*\prime}(b)}{B_H^*(b)} \cdot \mathcal{B}^*(0) + n \cdot \varphi_L^*(b) \cdot \frac{B_L^{*\prime}(b)}{B_L^*(b)} \cdot \mathcal{B}^*(b) \Big) \cdot \mathrm{d}b \\
\leq\;& \mathcal{B}^*(0) + \int_0^{\lambda^*} \Big(\frac{B_H^{*\prime}(b)}{B_H^*(b)} \cdot \mathcal{B}^*(0) + n \cdot \varphi_L^*(b) \cdot \frac{B_L^{*\prime}(b)}{B_L^*(b)} \cdot \mathcal{B}^*(b) \Big) \cdot \mathrm{d}b \\
=\;& \mathcal{B}^*(0) + \int_0^{\lambda^*} \mathcal{B}^{*\prime}(b) \cdot \mathrm{d}b - \int_0^{\lambda^*} n \cdot \big(1 - \varphi_L^*(b)\big) \cdot \frac{B_L^{*\prime}(b)}{B_L^*(b)} \cdot \mathcal{B}^*(b) \cdot \mathrm{d}b \\
=\;& 1 - \int_0^{\lambda^*} n \cdot \big(1 - \varphi_L(b)\big) \cdot \frac{1}{n \cdot (1 - b)} \cdot B_H^*(b) \cdot \frac{1 - \lambda^*}{1 - b} \cdot \mathrm{d}b \\
=\;& 1 - (1 - \lambda^*) \cdot \int_0^{\lambda^*} \big(1 - \varphi_L^*(b)\big) \cdot \frac{B_H^*(b)}{(1 - b)^2} \cdot \mathrm{d}b \\
=\;& 1 - (1 - \lambda^*) \cdot \int_1^2 \Big(\frac{n - (t-1)}{nt - (t-1)} \cdot t^2 \cdot e^{2-2t} \Big) \cdot \frac{t^2/4}{(t^2 \cdot e^{2-2t})^2} \cdot \Big(\frac{\mathrm{d}b}{\mathrm{d}t} \Big) \cdot \mathrm{d}t \\
=\;& 1 - (1 - \lambda^*) \cdot \int_1^2 \Big(\frac{n - (t-1)}{nt - (t-1)} \cdot t^2 \cdot e^{2-2t} \Big) \cdot \frac{t^2/4}{(t^2 \cdot e^{2-2t})^2} \cdot (2t^2 - 2t) \cdot e^{2-2t} \cdot \mathrm{d}t \\
=\;& 1 - (1 - \lambda^*) \cdot \int_1^2 \frac{n - (t-1)}{nt - (t-1)} \cdot \frac{t^2 - t}{2} \cdot \mathrm{d}t
\end{aligned}$$

$$\leq\ 1-(1-\lambda^*)\cdot\int_1^2\left(1-\frac{1}{n}\right)\cdot\frac{t-1}{2}\cdot \mathrm{d}t$$
$$=\ 1-\left(1-\frac{1}{n}\right)\cdot e^{-2}$$
$$\leq\ 1-(1-\varepsilon)\cdot e^{-2}.$$

The definition of $t\in[1,\,2]$ and the expressions of b, B_H^*, φ_L^* and $\frac{\mathrm{d}b}{\mathrm{d}t}$ in terms of t are given in the proof of Lemma 4.3. **Item 2** follows then. This finishes the proof of Lemma 4.3. ◻

4.2 Uniqueness of equilibria. In this part, we show that the focal equilibrium $\boldsymbol{s}^*=\{s_H^*\}\otimes\{s_L^*\}^{\otimes n}$ is the *unique* equilibrium for our modified instance $\boldsymbol{V}=V_H\otimes\{V_L\}^{\otimes n}$; thus the tight PoS bound is the same as the tight PoA bound. This uniqueness is the key ingredient of our PoS characterization. To have a better sense, let us explain the high-level ideas before giving the formal proof.

We first prove that bidders $\{L_i\}_{i\in[n]}$ must have identical strategies because they have identical value distributions (Lemma 4.5). Thus, we "truly" have just two kinds of bidders, H versus $\{L\}^{\otimes n}$. Then an equilibrium can be obtained by resolving an ordinal differential equation (ODE) in terms of the bid distributions for H and L (Lemma 4.8). Once the boundary conditions are specified, an ODE "usually" has one unique solution. Essentially, the possible non-uniqueness of equilibria stems from different boundary conditions.

The main technical part is to uniquely determine the boundary condition, namely every bidder H or L must have her bid support being exactly the interval $[0,\,\lambda^*]$. Compared with the tight PoA instances [JL22, Example 4], we modify bidder H's value distribution by putting a tiny probability mass at the zero value $\mathbf{Pr}[v_H=0]=\varepsilon$. In this way, every bidder H or L is enforced a zero bid, once this bidder has a zero value (Lemma 4.4). Then, it is easy to conclude that the identical bid support of bidders $\{L\}^{\otimes n}$ is exactly an interval $[0,\,\lambda]$ – having densities *almost everywhere* – since those bidders have an *uninterrupted* value support $\mathrm{supp}(V_L)=[0,\,1-\frac{2n-2}{2n-1}\cdot 2/e^2]$ (Lemma 4.6); but whether this supremum bid λ is exactly the $\lambda^*=1-4/e^2\approx 0.4587$ is still unclear.

However, determining the desirable boundary condition for bidder H is highly nontrivial. This bidder has an *interrupted* value support $\mathrm{supp}(V_H)=\{0,\,1\}$, so the above arguments fail to work. Instead, we first show that bidder H's bid support is the union $\{0\}\cup[\mu,\,\lambda]$ of (i) the zero bid $\{0\}$, which corresponds to the zero value $\mathbf{Pr}[v_H=0]=\varepsilon$; and (ii) an interval $[\mu,\,\lambda]$ for some $\mu\geq 0$ – having densities *almost everywhere* – which corresponds to the nonzero value $\mathbf{Pr}[v_H=1]=1-\varepsilon$.[4] The undesirable case $\mu>0$ is really possible, if we could slightly adjust Section 4, e.g., changing the "success $=\varepsilon$"/"failure $=1-\varepsilon$" probabilities of bidder H's Bernoulli random value. But under our particular construction, only the desirable case $\mu=0$ turns out to be possible; thus bidder H also has densities almost everywhere on the interval $[0,\,\lambda]$ (Lemma 4.7).

Provided with the desirable boundary conditions, we resolve the mentioned ODE, thus *uniquely* determining the bid support $[0,\,\lambda]=[0,\,\lambda^*]$ and the equilibrium – precisely the focal equilibrium $\boldsymbol{s}^*=\{s_H^*\}\otimes\{s_L^*\}^{\otimes n}$ (Lemma 4.8).

In the rest of Section 4.2, we start with a generic equilibrium $\boldsymbol{s}=\{s_H\}\otimes\{s_{L,\,i}\}_{i\in[n]}\in\mathbb{BNE}(\boldsymbol{V})$ and present the formal proof.

LEMMA 4.4. *Each bidder $\sigma\in\{H\}\cup\{L_i\}_{i\in[n]}$, on having a zero value $\{v_\sigma=0\}$, takes a zero bid $s_\sigma(v_\sigma)=0$ almost surely. Hence, bidder H takes a zero bid with probability $B_H(0)\geq V_H(0)=\varepsilon$ and each bidder L_i for $i\in[n]$ takes a zero bid with probability $B_{L,\,i}(0)\geq V_L(0)=\sqrt[n]{4/e^2}$. Further, the infimum bid $\gamma=\inf(\mathrm{supp}(\mathcal{B}))$ is zero $\gamma=0$.*

Proof. The all-zero value profile $\{\boldsymbol{v}=\boldsymbol{0}\}$ occurs with probability $V_H(0)\cdot(V_L(0))^n=\varepsilon\cdot 4/e^2>0$. Conditioned on this, the bid profile must also be all-zero $\{\boldsymbol{s}(\boldsymbol{v})=\boldsymbol{0}\}$, almost surely – Otherwise, with a nonzero probability, the allocated bidder $\mathrm{x}=\mathrm{x}(\boldsymbol{s}(\boldsymbol{v}))$ gains a *strictly* negative utility <0 since she has a nonzero bid $s_{\mathrm{x}}(v_{\mathrm{x}})>0$ and a zero value $v_{\mathrm{x}}=0$, which contradicts the equilibrium condition (Definition 2.1). We are considering Bayesian Nash Equilibria, namely the strategies $s_\sigma(v_\sigma)$ for $\sigma\in\{H\}\cup\{L_i\}_{i\in[n]}$ only depend on individual values v_σ. Therefore, for each individual bidder, a zero value $\{v_\sigma=0\}$ enforces a zero bid $\{s_\sigma(v_\sigma)=0\}$, almost surely. ◻

[4] A better interpretation is from the perspective of quantiles; then there is no ambiguity even in the case $\mu=0$.

Lemma 4.5. *Bidders* $\{L_i\}_{i\in[n]}$ *play identical strategies* $\{s_{L,i}\}_i = \{s_L\}^{\otimes n}$ *everywhere except on a zero-measure set of values. Hence, bid distributions* $\{B_{L,i}\}_{i\in[n]} = \{B_L\}^{\otimes n}$ *are identical and the bid-to-value mappings* $\{\varphi_{L,i}\}_{i\in[n]} = \{\varphi_L\}^{\otimes n}$ *are identical.*

Proof. Lemma 4.5 is almost a direct implication of [CH13, Corollary 3.2], which claims the same result for any subset of bidders $J = \{j_1, \ldots, j_k\} \subseteq [m]$ in an m-bidder First Price Auction that have *identical* value distributions $V_{j_1} \equiv \cdots \equiv V_{j_k}$. The only issue is that [CH13, Corollary 3.2] requires a tie-breaking rule $x_J \in \mathbb{FPA}$ that is *symmetric* for those bidders J. However, regarding our instance $\{H\} \cup \{L_i\}_{i\in[n]}$ and tie-breaking rule $x \in \mathbb{FPA}$ in Section 4:
(i) Tie-breaks at nonzero first-order bids $\{\max(\boldsymbol{s}(\boldsymbol{v})) > 0\}$ never occur, almost surely (Proposition 2.2).
(ii) Tie-breaks at a zero first-order bid $\{\max(\boldsymbol{s}(\boldsymbol{v})) = 0\}$, i.e., at the all-zero bid profile $\{\boldsymbol{s}(\boldsymbol{v}) = 0\}$, always favor bidder H and thus *is* symmetric for bidders $\{L_i\}_{i\in[n]}$.
Accordingly, the symmetry requirement on the tie-breaking rule fails just for a *zero-measure* set of values. Clearly, we can readopt the arguments for [CH13, Corollary 3.2] to derive Lemma 4.5. □

Recall that under the focal equilibrium $\boldsymbol{s}^*$, bidders $\{L_i\}_{i\in[n]}$ are non-monopoly bidders and have the same probability masses $B^*_L(0) = B_L(0) = \sqrt[n]{4/e^2}$. Below we show that this also holds for the considered equilibrium $\boldsymbol{s} = \{s_H\} \otimes \{s_L\}^{\otimes n}$.

Lemma 4.6. *Bidders* $\{L_i\}_{i\in[n]}$ *are non-monopoly bidders. Hence, the common bid distribution* B_L *has a probability mass* $B_L(0) = V_L(0) = \sqrt[n]{4/e^2}$ *at the zero bid and has densities almost everywhere over the bid support* $b \in [0, \lambda]$.

Proof. Bidders $\{L_i\}_{i\in[n]}$ have identical value/bid distributions $\{V_L\}^{\otimes n}$ and $\{B_L\}^{\otimes n}$. According to Definition 2.3, they either ALL are non-monopoly bidders $B_L(0) = V_L(0)$ or ALL are monopolists $B_L(0) > V_L(0)$. However, Proposition 2.4 (that there exists at most one monopolist) eliminates the second case. Hence, bidders $\{L_i\}_{i\in[n]}$ are non-monopoly bidders $B_L(0) = V_L(0)$.

For the sake of contradiction, assume that bid distribution B_L has no density around some bid $b \in [0, \lambda]$. Then, the competing bid distribution $(B_L(b))^n = \prod_{i\in[n]} B_{L,i}(b)$ for bidder H also has no density around this bid $b \in [0, \lambda]$. However, this contradicts Item 1 of Proposition 2.2. Refuting our assumption finishes the proof of Lemma 4.6. □

Lemma 4.7. *Bidder* H *is the (unique) monopolist. Hence, bid distribution* B_H *has a probability mass* $B_H(0) > V_H(0) = \varepsilon$ *at the zero bid and the bid-to-value mapping is constant* $\varphi_H(b) = 1$ *over the bid support* $b \in [0, \lambda]$.

Proof. Bidder H has (Section 4) a Bernoulli random value $\mathbf{Pr}[v_H = 0] = \varepsilon$ and $\mathbf{Pr}[v_H = 1] = 1 - \varepsilon$ and (Proposition 2.3) an increasing bid-to-value mapping $\varphi_H(b)$.

Assume to the contrary that bidder H is a non-monopoly bidder $B_H(0) = V_H(0) = \varepsilon$, namely a zero value $\{v_H = 0\}$ induces a zero bid $s_H(v_H) = 0$ almost surely, and a nonzero value $\{v_H = 1\}$ induces a nonzero bid $s_H(v_H) \in (0, \lambda]$ almost surely.

Consider the threshold bid $\mu := \inf\{b \in [0, \lambda] \mid \varphi_H(b) \geq 1\}$; this threshold bid $\mu \in [0, \lambda]$ is well defined regardless of our non-monopoly assumption for bidder H. **Fact 4.1** and **Fact 4.2** will be helpful for the later proof; only **Fact 4.2** relies on our non-monopoly assumption for bidder H.

Fact 4.1. **(I)** *Bid distribution* B_H *has no density on the interval* $b \in (0, \mu)$, *namely* $B_H(\mu) = B_H(0)$, *and has densities almost everywhere on the interval* $b \in (\mu, \lambda)$.
(II) *The bid-to-value mapping* $\varphi_H(b) \leq 1$ *for* $b \in [0, \lambda]$; *the equality holds when* $b \in [\mu, \lambda]$.

Proof. By the definition of $\mu = \inf\{b \in [0, \lambda] \mid \varphi_H(b) \geq 1\}$, a nonzero bid $s_H(v_H) \in (0, \lambda]$ due to the nonzero value $\{v_H = 1\}$ can be further restricted to the range $s_H(v_H) \in (\mu, \lambda]$. (Recall Item 2 of Proposition 2.2 that the bid CDF $B_H(b)$ is a *continuous* function over the bid support $b \in [0, \lambda]$.) Namely, bid distribution B_H has no density on the interval $b \in (0, \mu)$ and thus $B_H(\mu) = B_H(0)$.

For the sake of contradiction, assume that bid distribution B_H has no density around some bid $\beta \in (\mu, \lambda)$, namely $B'_H(\beta) = 0$. Then at this particular bid $\beta \in (\mu, \lambda)$, the two bid-to-value mappings $\varphi_H(b)$ and $\varphi_L(b)$ satisfy that

$$1 \;\leq\; \varphi_H(\beta) \;=\; \beta + \Big(n \cdot B'_L(\beta)/B_L(\beta)\Big)^{-1} \qquad \beta \in (\mu, \lambda)$$

$$\leq\ \varphi_L(\beta)\ =\ \beta + \Big((n-1)\cdot B_L'(\beta)/B_L(\beta) + B_H'(\beta)/B_H(\beta)\Big)^{-1}. \qquad\qquad B_H'(\beta)=0$$

This means value distribution V_L has densities around some value $v_\beta \geq 1$. Precisely, value distribution V_L can be reconstructed via the parametric equation $\{(v,\, V_L) = (\varphi_L(b),\, B_L(b)) \,|\, b \in [0,\, \lambda]\}$. Furthermore, bid distribution B_L has densities almost everywhere over the bid support $b \in [0,\, \lambda]$ (Lemma 4.6), including the particular bid $\beta \in (\mu,\, \lambda)$ for which $v_\beta = \varphi_L(\beta) \geq 1$. But this contradicts our construction – The hypothetical value $v_\beta = \varphi_L(\beta) \geq 1$ is bounded away from value distribution V_L's support $\mathrm{supp}(V_L) = [0,\, 1 - \frac{2n-2}{2n-1}\cdot 2/e^2]$ (Section 4).

Refuting the above assumption results in **Part (I)**: Bid distribution B_H has densities almost everywhere on $b \in (\mu,\, \lambda)$. All those densities stem from the value $\{v_H = 1\}$, since bidder H has a Bernoulli random value $v_H \in \{0,\, 1\}$ and (by assumption) is a non-monopoly bidder. Therefore, we have $\varphi_H(b) = 1$ for $b \in [\mu,\, \lambda]$. This together with monotonicity of the bid-to-value mapping $\varphi_H(b)$, immediately gives **Part (II)**. This finishes the proof. ∎

Fact 4.2. *Assume that bidder H is a non-monopoly bidder $B_H(0) = V_H(0) = \varepsilon$.*
(I) *The supremum bid $\lambda \leq \lambda^* = 1 - 4/e^2$.*
(II) $B_L(b) \geq B_L^*(b)$ *for* $b \in [0,\, \lambda]$.
(III) $B_L'(b)/B_L(b) \geq {B_L^*}'(b)/B_L^*(b)$ *for* $b \in [0,\, \lambda]$*; the equality holds when* $b \in [\mu,\, \lambda]$.
(IV) $B_H'(b)/B_H(b) \leq {B_H^*}'(b)/B_H^*(b)$ *for* $b \in [\mu,\, \lambda]$.

Proof. The bid-to-value mapping $\varphi_H(b) = b + \frac{1}{n}\cdot B_L(b)/B_L'(b) \leq 1$ over the bid support $b \in [0,\, \lambda]$; the equality holds when $b \in [\mu,\, \lambda]$. In contrast, the focal bid-to-value mapping $\varphi_H^*(b) = b + \frac{1}{n}\cdot B_L^*(b)/{B_L^*}'(b) = 1$ over the focal bid support $b \in [0,\, \lambda^*]$. Therefore, for $b \in [0,\, \min(\lambda,\, \lambda^*)]$ we have $B_L'(b)/B_L(b) \geq {B_L^*}'(b)/B_L^*(b)$ and thus

$$B_L(b)/B_L(0)\ =\ \exp\Big(\int_0^b B_L'(b)/B_L(b)\cdot \mathrm{d}x\Big)\ \geq\ B_L^*(b)/B_L^*(0)\ =\ \exp\Big(\int_0^b {B_L^*}'(b)/B_L^*(b)\cdot \mathrm{d}x\Big).$$

The two bid distributions have the same probability mass $B_L(0) = B_L^*(0) = V_L(0) = \sqrt[n]{4/e^2}$ at the zero bid, so we have $B_L(b) \geq B_L^*(b)$ for $b \in [0,\, \min(\lambda,\, \lambda^*)]$. To achieve the boundary conditions $B_L(\lambda) = 1$ and $B_L^*(\lambda^*) = 1$ at the respective supremum bids λ and λ^*, we must have **Part (I)** that $\lambda \leq \lambda^* = 1 - 4/e^2$.

Part (II) and **Part (III)**, including the equality $B_L'(b)/B_L(b) = {B_L^*}'(b)/B_L^*(b)$ for $b \in [\mu,\, \lambda]$, can be easily inferred from the above arguments.

Value distribution V_L can be reconstructed from EITHER bid distributions $B_H \otimes \{B_L\}^{\otimes n}$ OR the focal bid distributions $B_H^* \otimes \{B_L^*\}^{\otimes n}$, via the parametric equations $\{(v,\, V_L) = (\varphi_L(b),\, B_L(b)) \,|\, b \in [0,\, \lambda]\}$ or $\{(v,\, V_L) = (\varphi_L^*(b),\, B_L^*(b)) \,|\, b \in [0,\, \lambda^*]\}$. As Figure 2 suggests, this observation together with **Part (II)** implies that $\varphi_L(b) \geq \varphi_L^*(b)$ for $b \in [0,\, \lambda]$. Especially, on the restricted interval $b \in [\mu,\, \lambda]$, we can deduce that

$$\begin{aligned}\varphi_L(b)\ &=\ b + \Big((n-1)\cdot B_L'(b)/B_L(b) + B_H'(b)/B_H(b)\Big)^{-1}\\ &\geq\ \varphi_L^*(b)\ =\ b + \Big((n-1)\cdot {B_L^*}'(b)/B_L^*(b) + {B_H^*}'(b)/B_H^*(b)\Big)^{-1}.\end{aligned}$$

Rearranging this equation and applying **Part (III)** of **Fact 4.2** (that $B_L'(b)/B_L(b) = {B_L^*}'(b)/B_L^*(b)$ when $b \in [\mu,\, \lambda]$), we can conclude **Part (IV)** immediately. This finishes the proof of **Fact 4.2**. ∎

However, combining everything together, we can derive the following contradiction:

$$\begin{aligned}1\ =\ B_H(\lambda)\ &=\ B_H(\lambda)/B_H(\mu)\cdot B_H(\mu)\\ &=\ \exp\Big(\int_\mu^\lambda B_H'(b)/B_H(b)\cdot \mathrm{d}x\Big)\cdot B_H(\mu)\\ &=\ \exp\Big(\int_\mu^\lambda B_H'(b)/B_H(b)\cdot \mathrm{d}x\Big)\cdot \varepsilon && \textbf{Part (I) of Fact 4.1}\\ &\leq\ \exp\Big(\int_\mu^\lambda {B_H^*}'(b)/B_H^*(b)\cdot \mathrm{d}x\Big)\cdot \varepsilon && \textbf{Part (IV) of Fact 4.2}\end{aligned}$$

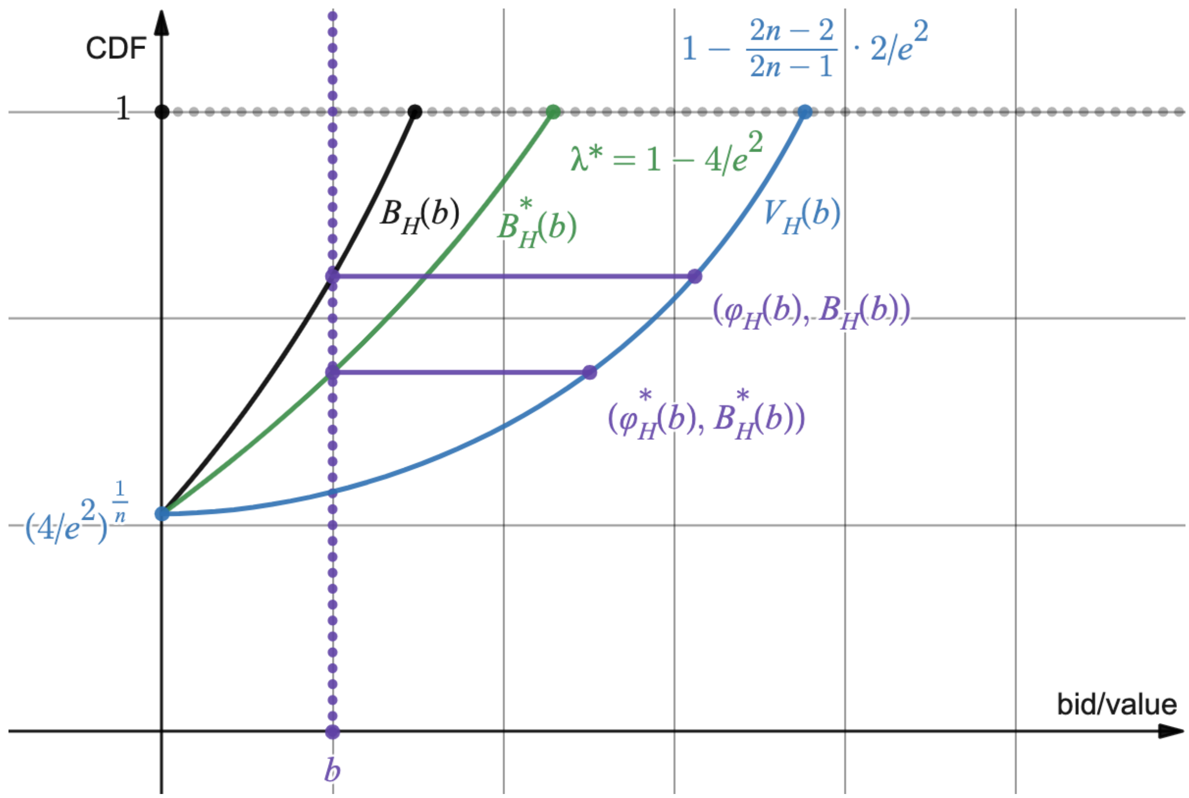

Figure 2: Demonstration for the proof of **Part (IV)** of **Fact 4.2**.

$$= B_H^*(\lambda)/B_H^*(\mu) \cdot \varepsilon$$
$$\leq 4\varepsilon < 1/2.$$

Here the last line uses $B_H^*(\mu) \geq B_H^*(0) = 1/4$, $B_H^*(\lambda) \leq B_H^*(\lambda^*) = 1$, and $\varepsilon \in (0, 1/8)$; all of which can be found from Section 4 and **Part (I)** of **Fact 4.2**.

Refute our assumption: Bidder H is the unique monopolist $B_H(0) > V_H(0) = \varepsilon$; the probability mass $B_H(0) > \varepsilon$ at the zero bid stems from BOTH a zero value $\{v_H = 0\}$ (by the whole amount $\varepsilon = \mathbf{Pr}[v_H = 0]$) AND a nonzero value $\{v_H = 1\}$ (by a partial amount $B_H(0) - \varepsilon \leq \mathbf{Pr}[v_H = 1]$). This implies that the threshold bid $\mu = \inf\{b \in [0, \lambda] \mid \varphi_H(b) \geq 1\}$ is zero $\mu = 0$. As a consequence, we can infer Lemma 4.7 from **Fact 4.1**. This finishes the proof. ∎

Lemma 4.8. (Uniqueness of Equilibria) *The following hold:*

1. *The supremum bid $\lambda = 1 - 4/e^2$, the same as the focal supremum bid $\lambda = \lambda^*$.*
2. *Bid distributions $\{B_{L,i}\}_{i \in [n]} = \{B_L\}^{\otimes n}$ are given by $B_L(b) = \sqrt[n]{(1-\lambda)/(1-b)}$ for $b \in [0, \lambda]$, the same as the focal bid distribution $B_L \equiv B_L^*$.*
3. *Bid distribution B_H is given by the implicit equation $b = 1 - 4B_H \cdot e^{2-4\sqrt{B_H}}$ for $B_H \in [1/4, 1]$, the same as the focal bid distribution $B_H \equiv B_H^*$.*

Proof. Following Lemma 4.7, over the bid support $b \in [0, \lambda]$, bidder H has a constant bid-to-value mapping $\varphi_H(b) = b + \frac{1}{n} \cdot B_L(b)/B_L'(b) = 1$. By resolving this ODE, under the boundary condition $B_L(0) = \sqrt[n]{4/e^2}$ at the infimum bid $= 0$, we have $B_L(b) = \sqrt[n]{(4/e^2)/(1-b)}$ for $b \in [0, \lambda]$. Plugging this CDF formula into the other boundary condition $B_L(\lambda) = 1$ at the supremum bid $= \lambda$, we can deduce that $\lambda = 1 - 4/e^2$. **Item 1** and **Item 2** follow then.

It remains to show **Item 3**. By construction (Section 4), bidders $\{L_i\}_{i\in[n]}$ have identical value distributions $\{V_L\}^{\otimes n}$ given by $V_L(1-\frac{n-(t-1)}{nt-(t-1)}\cdot t^2\cdot e^{2-2t}) = \sqrt[n]{4/t^2\cdot e^{2t-4}}$ for $t\in[1,\,2]$. Those value distributions also can be reconstructed through the parametric equation $\big\{(v,\,V_L)=(\varphi_L(b),\,B_L(b))\,\big|\,b\in[0,\,\lambda]\big\}$. As a combination, in terms of $t\in[1,\,2]$, the bid-to-value mapping φ_L can be rewritten as follows:

$$\varphi_L = 1-\frac{n-(t-1)}{nt-(t-1)}\cdot t^2\cdot e^{2-2t}. \tag{4.1}$$

Similarly, we can deduce that $\sqrt[n]{4/t^2\cdot e^{2t-4}} = B_L(b) = \sqrt[n]{(1-\lambda)/(1-b)}$ and thus rewrite the bid $b = 1-t^2\cdot e^{2-2t}$ for $t\in[1,\,2]$. Then, the derivative $\mathrm{d}b/\mathrm{d}t = (2t^2-2t)\cdot e^{2-2t}$.

On the other hand, each bidder L_i for $i\in[n]$ competes with bidders $\{H\}\cup\{L_j\}_{j\in[n]\setminus\{i\}}$, hence the competing bid distribution $\mathcal{B}_{-L}(b) = B_H(b)\cdot\big(B_L(b)\big)^{n-1}$. Consequently, in terms of $t\in[1,\,2]$, the bid-to-value mapping $\varphi_L = b+\frac{\mathcal{B}_{-L}}{\mathrm{d}\mathcal{B}_{-L}/\mathrm{d}b}$ also can be rewritten as follows:

$$\begin{aligned}
\varphi_L &= b+\Big((n-1)\cdot\frac{\mathrm{d}B_L/\mathrm{d}b}{B_L}+\frac{\mathrm{d}B_H/\mathrm{d}b}{B_H}\Big)^{-1}\\
&= b+\Big(\frac{n-1}{n}\cdot\frac{1}{1-b}+\frac{\mathrm{d}B_H/\mathrm{d}t}{B_H}\cdot\frac{1}{(2t^2-2t)\cdot e^{2-2t}}\Big)^{-1}\\
&= (1-t^2\cdot e^{2-2t})+\Big(\frac{n-1}{n}\cdot\frac{1}{t^2\cdot e^{2-2t}}+\frac{\mathrm{d}B_H/\mathrm{d}t}{B_H}\cdot\frac{1}{(2t^2-2t)\cdot e^{2-2t}}\Big)^{-1}.
\end{aligned} \tag{4.2}$$

Here the second line applies $B_L(b) = \sqrt[n]{(1-\lambda)/(1-b)}$ and $\mathrm{d}b/\mathrm{d}t = (2t^2-2t)\cdot e^{2-2t}$, and the last line applies $b = 1-t^2\cdot e^{2-2t}$.

The above two formulas for the φ_L must be identical for $t\in[1,\,2]$, so we can deduce that

$$\begin{aligned}
&\text{Equation (4.1)} \;=\; \text{Equation (4.2)}\\
\iff\;& \Big(1-\frac{n-(t-1)}{nt-(t-1)}\Big)\cdot t^2\cdot e^{2-2t} \;=\; \Big(\frac{n-1}{n}\cdot\frac{1}{t^2\cdot e^{2-2t}}+\frac{\mathrm{d}B_H/\mathrm{d}t}{B_H}\cdot\frac{1}{(2t^2-2t)\cdot e^{2-2t}}\Big)^{-1}\\
\iff\;& \frac{nt-(t-1)}{n\cdot(t-1)}\cdot\frac{1}{t^2\cdot e^{2-2t}} \;=\; \frac{n-1}{n}\cdot\frac{1}{t^2\cdot e^{2-2t}}+\frac{\mathrm{d}B_H/\mathrm{d}t}{B_H}\cdot\frac{1}{(2t^2-2t)\cdot e^{2-2t}}\\
\iff\;& \frac{1}{t-1}\cdot\frac{1}{t^2\cdot e^{2-2t}} \;=\; \frac{\mathrm{d}B_H/\mathrm{d}t}{B_H}\cdot\frac{1}{(2t^2-2t)\cdot e^{2-2t}}\\
\iff\;& \frac{2}{t} \;=\; \frac{\mathrm{d}B_H/\mathrm{d}t}{B_H}\\
\iff\;& \frac{\mathrm{d}}{\mathrm{d}t}(2\ln t) \;=\; \frac{\mathrm{d}}{\mathrm{d}t}(\ln B_H).
\end{aligned} \tag{4.3}$$

Especially, when $t=2$, the bid $b = 1-t^2\cdot e^{2-2t} = 1-4/e^2$ achieves the supremum bid $\lambda = 1-4/e^2$ and we have the boundary condition $B_H(b) = B_H(\lambda) = 1$. Resolving ODE (4.3) under this boundary condition, we derive that $2\ln(t/2) = \ln B_H$ and thus $B_H = t^2/4$. Plugging this into the formula $b = 1-t^2\cdot e^{2-2t}$ for $t\in[1,\,2]$ gives the implicit equation $b = 1-4B_H\cdot e^{2-4\sqrt{B_H}}$ for $B_H\in[1/4,\,1]$, as desired. This finishes the proof of Lemma 4.8. ∎

5 Bayesian Nash Equilibria for Correlated Valuations.

In this section, we give the tight PoS bound for correlated valuations.

Theorem 5.1. (Tight PoS) *Regarding Bayesian Nash Equilibria for correlated valuations, the Price of Stability is $1-1/e\approx 0.6321$.*

The proof framework is similar to that in the previous section, given that this tight PoS bound $=1-1/e$ also coincides with the tight PoA bound by [ST13, Syr14]. We slightly modify the tight PoA instance from [Syr14, Appendix A.2] into the next instance and prove that it has one unique equilibrium. (Namely, by setting $\varepsilon = 0$, Section 5 is precisely the original tight PoA instance.)

Example. Given an arbitrarily small constant $\varepsilon \in (0,\, 1)$, consider the instance $\{H,\, L_1,\, L_2\}$ in terms of the joint value distribution $\boldsymbol{v} = (v_H,\, v_{L,\,1},\, v_{L,\,2}) \sim \boldsymbol{V}$.

- Bidder H has an *independent* Bernoulli random value $\mathbf{Pr}[v_H = 0] = \varepsilon$ and $\mathbf{Pr}[v_H = 1] = 1 - \varepsilon$. Denote by V_H this (marginal) value distribution.
- Bidders L_1 and L_2 have *perfectly correlated* and *identical* values $v_{L,\,1} \equiv v_{L,\,2}$, which follow the (marginal) value distribution $V_L(v) = \frac{\varepsilon + 1/e}{\varepsilon + 1 - v}$ for $v \in [0,\, 1 - 1/e]$.

The considered First Price Auction $\mathrm{x} \in \mathbb{FPA}$, under the all-zero bid profile $\boldsymbol{b} = (b_H,\, b_{L,\,1},\, b_{L,\,2}) = \mathbf{0}$, favors bidder H, but otherwise is arbitrary. ∎

5.1 The focal equilibrium.

The ***focal*** *strategy profile* $\boldsymbol{s}^* = \{s_H^*\} \otimes \{s_L^*\}^{\otimes 2}$ is given as follows.

- Bidder H has a *fixed* strategy $s_H^*(v_H) \equiv 0$ for $v_H \in \{0,\, 1\}$.
- Bidders L_1 and L_2 have identical and *truthful* strategies $s_L^*(v) \equiv v$ for $v \in [0,\, 1 - 1/e]$.

LEMMA 5.1. (EQUILIBRIUM) *The focal strategy profile $\boldsymbol{s}^* = s_H^* \otimes \{s_L^*\}^{\otimes 2}$ forms a Bayesian Nash Equilibrium $\boldsymbol{s}^* \in \mathbb{BNE}(\boldsymbol{V})$.*

Proof. Bidder L_1 satisfies the equilibrium condition: Given a specific bid $s_L^*(v_1) = b \in [0,\, 1 - 1/e]$, bidder L_1' value and bidder L_2's value/bid ALL must be the same $v_1 = v_2 = s_L^*(v_1) = s_L^*(v_2) = b$. Clearly, bidder L_1 gains a zero utility $= 0$ and cannot benefit from a deviation bid $b' \neq b$, namely a higher bid $b' > b$ gives a nonpositive utility ≤ 0 and a lower bid $b' < b = s_L^*(v_2)$ makes bidder L_1 lose to bidder L_2. By symmetry, bidder L_2 also meets the equilibrium condition.

Bidder H's competing bid distribution $\max(s_L^*(v_1),\, s_L^*(v_2)) \sim \mathcal{B}_{-H}^*$ is exactly bidders L_1 and L_2's value distribution V_L. For any value $v_H \in \{0,\, 1\}$ and any bid $b \in [0,\, 1 - 1/e]$, bidder H gains an interim utility $= (v_H - b) \cdot V_L(b) = \frac{v_H - b}{\varepsilon + 1 - b} \cdot (\varepsilon + 1/e)$. We can easily verify that, under the either value $v_H \in \{0,\, 1\}$, the zero bid $b = 0$ is the *unique* maximizer for this interim utility formula; thus bidder H also satisfies the equilibrium condition. This finishes the proof. ∎

LEMMA 5.2. (EFFICIENCY) *The following hold for the focal equilibrium $\boldsymbol{s}^* = \{s_H^*\} \otimes \{s_L^*\}^{\otimes 2}$:*

1. *The expected optimal Social Welfare $\mathsf{OPT}(\boldsymbol{V},\, \boldsymbol{s}^*) \geq 1 - \varepsilon$.*

2. *The expected auction Social Welfare $\mathsf{FPA}(\boldsymbol{V},\, \boldsymbol{s}^*) \leq 1 - 1/e$.*

Proof. Let us prove **Items 1 and 2** one by one.

Item 1. The realized optimal Social Welfare is at least bidder H's realized value $v_H \sim V_H$, namely a Bernoulli random value $\mathbf{Pr}[v_H = 0] = \varepsilon$ and $\mathbf{Pr}[v_H = 1] = 1 - \varepsilon$ (Section 5). In expectation, we have $\mathsf{OPT}(\boldsymbol{V},\, \boldsymbol{s}^*) \geq \mathbf{E}[v_H] = 1 - \varepsilon$. **Item 1** follows then.

Item 2. The first-order bid distribution $\max(\boldsymbol{s}^*(\boldsymbol{v})) \sim \mathcal{B}^*$ is exactly bidders L_1 and L_2's value distribution V_L. Conditioned on a zero first-order bid $\{\max(\boldsymbol{s}^*(\boldsymbol{v})) = 0\}$, bidder H wins and the realized auction Social Welfare is her value $v_H \sim V_H$. And conditioned on a nonzero first-order bid $\{\max(\boldsymbol{s}^*(\boldsymbol{v})) > 0\}$, either bidder L_1 or bidder L_2 wins and the realized auction Social Welfare is their identical value $v_1 = v_2 = s_L^*(v_1) = s_L^*(v_2) = \max(\boldsymbol{s}^*(\boldsymbol{v})) > 0$. In expectation, we have

$$\begin{aligned}
\mathsf{FPA}(\boldsymbol{V},\, \boldsymbol{s}^*) &= \mathbf{E}[V_H] \cdot V_L(0) + \mathbf{E}[V_L] \\
&= (1 - \varepsilon) \cdot \frac{\varepsilon + 1/e}{\varepsilon + 1} + \int_0^{1 - 1/e} \big(1 - V_L(v)\big) \cdot \mathrm{d}v \\
&= 1 - 1/e - (\varepsilon + 1/e) \cdot \Big(\frac{2\varepsilon}{1 + \varepsilon} - \ln\Big(\frac{e\varepsilon + 1}{\varepsilon + 1}\Big)\Big) \\
&\leq 1 - 1/e.
\end{aligned}$$

Here the last two steps can be easily verified via elementary algebra. This finishes the proof. ∎

5.2 Uniqueness of equilibria. In this part, we prove the focal equilibrium $\boldsymbol{s}^*$ is the unique equilibrium for Section 5. Once again, we start with a generic equilibrium $\boldsymbol{s} = \{s_H,\, s_{L,\,1},\, s_{L,\,2}\} \in \mathbb{BNE}(\boldsymbol{V})$.

LEMMA 5.3. *Each bidder $i \in \{H,\, L_1,\, L_2\}$ cannot overbid, $s_i(v) \le v$ almost surely over the randomness of the strategy s_i, everywhere $v \in \mathrm{supp}(V_i)$ except on a zero-measure set of values.*

Proof. First, on having a zero value $\{v_i = 0\}$, each bidder $i \in \{H,\, L_1,\, L_2\}$ has a zero bid $s_i(v_i) = 0$ almost surely. Otherwise, with a nonzero probability > 0, the following event occurs.
$\{\boldsymbol{v} = \boldsymbol{0} \wedge \boldsymbol{s}(\boldsymbol{v}) \neq \boldsymbol{0}\}$: The value profile $\boldsymbol{v}$ is all-zero but the bid profile $\boldsymbol{s}(\boldsymbol{v})$ is not.
But conditioned on this, the allocated bidder $\mathrm{x}(\boldsymbol{s}(\boldsymbol{v}))$ realizes a negative utility $= -\max(\boldsymbol{s}(\boldsymbol{v})) < 0$, which contradicts the equilibrium condition.

Second, on having a nonzero value $\{v_H = 1\}$, bidder H cannot overbid $s_H(v_H) \le v_H$. (Recall that value $v_H \sim V_H$ is independent from the other two values $v_{L,\,1} \equiv v_{L,\,2} \sim V_L$.) Otherwise, with a nonzero probability > 0, the next two events occur simultaneously.
$\{v_H = 1 \wedge s_H(v_H) > 1\}$: Bidder H has a nonzero value and overbids;
$\{v_{L,\,1} = v_{L,\,2} = s_{L,\,1}(v_{L,\,1}) = s_{L,\,1}(v_{L,\,1}) = 0\}$: Bidders L_1 and L_2 have zero values and zero bids.
But conditioned on this, bidder H gets allocated and realizes a negative utility $= v_H - s_H(v_H) < 0$, which contradicts the equilibrium condition.

Third, on having nonzero values $(v_{L,\,1} \equiv v_{L,\,2}) > 0$, bidders L_1 and L_2 cannot overbid. Otherwise, with a nonzero probability > 0, the next two events occur simultaneously.
$\{\max(s_{L,\,1}(v_{L,\,1}),\, s_{L,\,2}(v_{L,\,2})) > (v_{L,\,1} \equiv v_{L,\,2}) > 0\}$: Bidders L_1 and L_2 have identical and nonzero values > 0 and at least one of them overbids;
$\{v_H = s_H(v_H) = 0\}$: Bidder H has a zero value and a zero bid.
But conditioned on this, the allocated bidder $\mathrm{x}(\boldsymbol{s}(\boldsymbol{v})) \in \mathrm{argmax}(\boldsymbol{s}(\boldsymbol{v}))$ is the higher bidder between L_1 and L_2, realizing a negative utility < 0, which contradicts the equilibrium condition.

This finishes the proof of Lemma 5.3. $\square$

LEMMA 5.4. *Bidders L_1 and L_2 play the truthful strategies, $s_{L,\,1}(v) = s_{L,\,2}(v) = v$ almost surely, everywhere $v \in [0,\, 1 - 1/e]$ except on a zero-measure set of values.*

Proof. Following Lemma 5.3, bidders L_1 and L_2 cannot overbid $s_{L,\,1}(v),\, s_{L,\,2}(v) \not> v$ and play the truthful strategies $s_{L,\,1}(0) = v_{L,\,2}(0) = 0$ on having the zero values. It remains to show that bidders L_1 and L_2 also cannot "shade" their bids, namely $s_{L,\,1}(v),\, s_{L,\,2}(v) \not< v$ for $v > 0$.

Assume the opposite: For some nonzero value $(v_{L,\,1} \equiv v_{L,\,2}) = v > 0$, either or both of $\{L_1,\, L_2\}$ "shades" her bid with a nonzero probability $\mathbf{Pr}_{s_{L,\,1},\, s_{L,\,2}}[\min(s_{L,\,1}(v),\, s_{L,\,2}(v)) < v] > 0$. Let us do case analysis conditioned on the event $\mathcal{E} = \{(v_{L,\,1} \equiv v_{L,\,2}) = v > 0\}$:

- **Case (I).** Exactly one bidder $\in \{L_1,\, L_2\}$ "shades" her bid with a nonzero probability.

  Without loss of generality, bidder L_1 plays the shade strategy $\mathbf{Pr}_{s_{L,\,1}}[s_{L,\,1}(v) < v] > 0$, while bidder L_2 plays the truthful strategy $\mathbf{Pr}_{s_{L,\,2}}[s_{L,\,2}(v) = v] = 1$. But if so, bidder L_2 can benefit from a certain deviation bid b^* against the current zero utility $= 0$ from the *truthful* strategy. Specifically, bidder L_1's infimum strategy $\underline{s}_1 = \inf(\mathrm{supp}(s_{L,\,1}(v)))$ is bounded away from the considered value $\underline{s}_1 < v$. Using the deviation bid $b^* = \frac{1}{2}(\underline{s}_1 + v)$, bidder L_2 realizes a positive utility $= v - b^* > 0$ on winning, and wins with a nonzero probability > 0: Independently,
  (i) bidder H loses with probability $\ge \mathbf{Pr}_{v_H,\, s_H}[s_H(v_H) < b^*] \ge \mathbf{Pr}_{v_H,\, s_H}[v_H = 0] = \varepsilon$, because bidder H on having a zero value $\{v_H = 0\}$ also has a zero bid $s_H(v_H) = 0$ (Lemma 5.3);
  (ii) bidder L_1 loses with a nonzero probability $\ge \mathbf{Pr}_{s_{L,\,1}}[s_{L,\,1}(v) < b^*] > 0$, as a consequence of $b^* = \frac{1}{2}(\underline{s}_1 + v) > \underline{s}_1 = \inf(\mathrm{supp}(s_{L,\,1}(v)))$.

  Thus, bidder L_2 can benefit from a certain deviation bid b^*. This contradicts the equilibrium condition and refutes **Case (I)**.

- **Case (II).** Each bidder $\in \{L_1,\, L_2\}$ "shades" her bid with a nonzero probability.

  Consider the infimum strategies $\underline{s}_1 = \inf(\mathrm{supp}(s_{L,\,1}(v))) < v$ and $\underline{s}_2 = \inf(\mathrm{supp}(s_{L,\,2}(v))) < v$ and their likelihoods $p_1 = \mathbf{Pr}_{s_{L,\,1}}[s_{L,\,1}(v) = \underline{s}_1]$ and $p_2 = \mathbf{Pr}_{s_{L,\,2}}[s_{L,\,2}(v) = \underline{s}_2]$.

– **Case (a).** $\underline{s}_1 = \underline{s}_2 = \underline{s}$ for some bid $\underline{s} \in [0, v]$ and $p_1, p_2 > 0$.

A tiebreak at the bid $\underline{s}$ occurs with a nonzero probability $\mathbf{Pr}_{v_H, \boldsymbol{s}}[\max(\boldsymbol{s}(\boldsymbol{v})) = \underline{s} \,|\, \mathcal{E}] = \mathbf{Pr}_{v_H, s_H}[s_H(v_H) \le \underline{s}] \cdot p_1 \cdot p_2 > 0$, since bidder H on having a zero value $\mathbf{Pr}[v_H = 0] = \varepsilon$ also has a zero bid $s_H(v_H) = 0$ (Lemma 5.3). In this tiebreak $\mathcal{E} \wedge \{\max(\boldsymbol{s}(\boldsymbol{v})) = \underline{s}\}$, at least one bidder between L_1 and L_2 loses with a nonzero probability > 0, say bidder L_1. But this means, using a higher but close enough deviation bid $b^* \searrow \underline{s}$, bidder L_1 realizes an arbitrarily close positive utility $= (v - b^*) \nearrow (v - \underline{s}) > 0$ on winning, yet the winning probability increases by a nonzero amount $\ge \mathbf{Pr}_{v_H, \boldsymbol{s}}[\max(\boldsymbol{s}(\boldsymbol{v})) = \underline{s} \,|\, \mathcal{E}] > 0$.

Thus, bidder L_1 can benefit from a higher but close enough deviation bid $b^* \searrow \underline{s}$. This contradicts the equilibrium condition and refutes **Case (a)**.

– **Case (b).** Either $\underline{s}_1 \ne \underline{s}_2$ or $p_1 \cdot p_2 = 0$.

If $\underline{s}_1 \ne \underline{s}_2$, without loss of generality we have $\underline{s}_1 < \underline{s}_2 < v$. But if so, bidder L_1 gains a nonzero utility > 0 from any bid $b^* \in (\underline{s}_2, v)$, in contrast to a zero utility $= 0$ from any bid $\in [\underline{s}_1, \underline{s}_2)$. The current strategy $s_{L,1}(v)$ has densities on the "useless" bids $\in [\underline{s}_1, \underline{s}_2)$ and cannot be utility-optimal.

If $\underline{s}_1 = \underline{s}_2 = \underline{s}$ for some bid $\underline{s} \in [0, v]$ and $p_1 \cdot p_2 = 0$, without loss of generality we have $p_1 = 0$. But if so, bidder L_2 gains a nonzero utility > 0 from any bid $b^* \in [\underline{s}, v]$ that is bounded away from both $\underline{s}$ and v, in contrast to an arbitrarily small utility $\searrow 0$ (as the winning probability $\searrow 0$) when her bid $\in [\underline{s}, v]$ approaches the infimum strategy $\searrow \underline{s}$. The current strategy $s_{L,2}(v)$ has densities in any neighborhood around the infimum strategy $\underline{s}$ and cannot be utility-optimal.

Thus, at least one bidder between L_1 and L_2 can benefit from a certain deviation bid b^*. This contradicts the equilibrium condition and refutes **Case (b)**.

In sum, **Case (II)** gets refuted.

Refute our assumption: Bidders L_1 and L_2 play truthful strategies $s_{L,1}(v) \equiv s_{L,2}(v) \equiv v$ everywhere $v \in [0, 1 - 1/e]$. This finishes the proof. ∎

LEMMA 5.5. *Bidder H has a fixed strategy $s_H(v_H) \equiv 0$ for $v_H \in \{0, 1\}$.*

Proof. We reuse the arguments for Lemma 5.1. Given a value $v_H \in \{0, 1\}$ and a bid $b \in [0, 1 - 1/e]$, bidder H gains an interim utility $= (v_H - b) \cdot V_L(b) = \frac{v_H - b}{\varepsilon + 1 - b} \cdot (\varepsilon + 1/e)$. The zero bid $b = 0$, under the either value $v_H \in \{0, 1\}$, is always the UNIQUE maximizer for this interim utility formula. ∎

Acknowledgements.

We are grateful to Xi Chen and Tao Lin for invaluable discussions.

Y.J. is supported by NSF grants IIS-1838154, CCF-1563155, CCF-1703925, CCF-1814873, CCF-2106429, and CCF-2107187. P.L. is supported by Science and Technology Innovation 2030 – "New Generation of Artificial Intelligence" Major Project No.(2018AAA0100903), NSFC grant 61922052 and 61932002, Innovation Program of Shanghai Municipal Education Commission, Program for Innovative Research Team of Shanghai University of Finance and Economics, and Fundamental Research Funds for Central Universities.

References

[ADK+08] Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM J. Comput.*, 38(4):1602–1623, 2008.

[AHN+19] Saeed Alaei, Jason D. Hartline, Rad Niazadeh, Emmanouil Pountourakis, and Yang Yuan. Optimal auctions vs. anonymous pricing. *Games and Economic Behavior*, 118:494–510, 2019.

[CH13] Shuchi Chawla and Jason D. Hartline. Auctions with unique equilibria. In Michael J. Kearns, R. Preston McAfee, and Éva Tardos, editors, *Proceedings of the fourteenth ACM Conference on Electronic Commerce, EC 2013, Philadelphia, PA, USA, June 16-20, 2013*, pages 181–196. ACM, 2013.

[CKK+15] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, Maria Kyropoulou, Brendan Lucier, Renato Paes Leme, and Éva Tardos. Bounding the inefficiency of outcomes in generalized second price auctions. *J. Econ. Theory*, 156:343–388, 2015.

[CP14] Yang Cai and Christos H. Papadimitriou. Simultaneous bayesian auctions and computational complexity. In Moshe Babaioff, Vincent Conitzer, and David A. Easley, editors, *ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014*, pages 895–910. ACM, 2014.

[FGH+21] Aris Filos-Ratsikas, Yiannis Giannakopoulos, Alexandros Hollender, Philip Lazos, and Diogo Poças. On the complexity of equilibrium computation in first-price auctions. In Péter Biró, Shuchi Chawla, and Federico Echenique, editors, *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, pages 454–476. ACM, 2021.

[HHT14] Jason D. Hartline, Darrell Hoy, and Sam Taggart. Price of anarchy for auction revenue. In Moshe Babaioff, Vincent Conitzer, and David A. Easley, editors, *ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014*, pages 693–710. ACM, 2014.

[HKMN11] Avinatan Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. Non-price equilibria in markets of discrete goods. In Yoav Shoham, Yan Chen, and Tim Roughgarden, editors, *Proceedings 12th ACM Conference on Electronic Commerce (EC-2011), San Jose, CA, USA, June 5-9, 2011*, pages 295–296. ACM, 2011.

[HTW18] Darrell Hoy, Samuel Taggart, and Zihe Wang. A tighter welfare guarantee for first-price auctions. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 132–137. ACM, 2018.

[JL22] Yaonan Jin and Pinyan Lu. First-price auction is $1 - 1/e^2$ efficient. To appear in *Proceedings of the 63rd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 2022.

[JLQ+19] Yaonan Jin, Pinyan Lu, Qi Qi, Zhihao Gavin Tang, and Tao Xiao. Tight approximation ratio of anonymous pricing. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019.*, pages 674–685, 2019.

[JSSZ02] Matthew O Jackson, Leo K Simon, Jeroen M Swinkels, and William R Zame. Communication and equilibrium in discontinuous games of incomplete information. *Econometrica*, 70(5):1711–1740, 2002.

[KP99] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In Christoph Meinel and Sophie Tison, editors, *STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science, Trier, Germany, March 4-6, 1999, Proceedings*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413. Springer, 1999.

[KZ12] Todd R Kaplan and Shmuel Zamir. Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory*, 50(2):269–302, 2012.

[Leb96] Bernard Lebrun. Existence of an equilibrium in first price auctions. *Economic Theory*, 7(3):421–443, 1996.

[Leb99] Bernard Lebrun. First price auctions in the asymmetric n bidder case. *International Economic Review*, 40(1):125–142, 1999.

[Leb06] Bernard Lebrun. Uniqueness of the equilibrium in first-price auctions. *Games Econ. Behav.*, 55(1):131–151, 2006.

[MR00a] Eric Maskin and John Riley. Asymmetric auctions. *The review of economic studies*, 67(3):413–438, 2000.

[MR00b] Eric Maskin and John Riley. Equilibrium in sealed high bid auctions. *The Review of Economic Studies*, 67(3):439–454, 2000.

[MR03] Eric Maskin and John Riley. Uniqueness of equilibrium in sealed high-bid auctions. *Games Econ. Behav.*, 45(2):395–409, 2003.

[Plu92] Michael Plum. Characterization and computation of nash-equilibria for auctions with incomplete information. *International Journal of Game Theory*, 20(4):393–418, 1992.

[Rou15] Tim Roughgarden. Intrinsic robustness of the price of anarchy. *J. ACM*, 62(5):32:1–32:42, 2015.

[ST13] Vasilis Syrgkanis and Éva Tardos. Composable and efficient mechanisms. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 211–220. ACM, 2013.

[Syr14] Vasileios Syrgkanis. *Efficiency of mechanisms in complex markets.* Cornell University, 2014.

[SZ90] Leo K Simon and William R Zame. Discontinuous games and endogenous sharing rules. *Econometrica: Journal of the Econometric Society*, pages 861–872, 1990.

[Vic61] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

[Vil09] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.

A Bayesian (Coarse) Correlated Equilibria.

We first introduce the definition of joint strategies and Bayesian (Coarse) Correlated Equilibria.

Definition A.1. (Joint Strategies) *A joint strategy profile $\boldsymbol{s}(\boldsymbol{v}) \equiv (s_i(\boldsymbol{v}))_{i\in[n]}$ involves $n \geq 1$ many n-variate functions; each one maps the whole value profile $\boldsymbol{v} \in \mathbb{R}^n_{\geq 0}$ to a nonnegative bid $s_i(\boldsymbol{v}) \geq 0$. (When the functions $(s_i(\boldsymbol{v}))_{i\in[n]}$ degenerate into* univariates *$s_i(\boldsymbol{v}) \equiv s_i(v_i)$, the strategy profile $\boldsymbol{s}(\boldsymbol{v})$ degenerates into an* independent *strategy profile, as before for* Bayesian Nash Equilibria*.)*

DEFINITION A.2. (Bayesian (Coarse) Correlated Equilibria) *Given a joint value distribution* $\boldsymbol{V} \in \mathbb{V}_{\mathsf{joint}}$*, a tie-breaking rule* $\mathrm{x} \in \mathbb{FPA}$*, and a precision* $\delta > 0$*:*

- *A* δ*-approximate* Bayesian Correlated Equilibrium $\boldsymbol{s} \in \mathbb{BCE}(\boldsymbol{V}, \mathrm{x}, \delta)$ *is a joint strategy profile that, for any bidder* $i \in [n]$*, value* $v \in \mathrm{supp}_i(\boldsymbol{v})$*, bid* $b \in \mathrm{supp}_i(\boldsymbol{s}(\boldsymbol{v}))$*, and deviation bid* $b^* \geq 0$*,*

$$\mathop{\mathbf{E}}_{\boldsymbol{v},\, \boldsymbol{s},\, \mathrm{x}}\big[\, u_i(v_i,\, \boldsymbol{s}(\boldsymbol{v})) \,\big|\, v_i = v,\, s_i(\boldsymbol{v}) = b \,\big] \;\geq\; \mathop{\mathbf{E}}_{\boldsymbol{v},\, \boldsymbol{s},\, \mathrm{x}}\big[\, u_i(v_i,\, b^*,\, \boldsymbol{s}_{-i}(\boldsymbol{v})) \,\big|\, v_i = v,\, s_i(\boldsymbol{v}) = b \,\big] - \delta.$$

- *A* δ*-approximate* Bayesian Coarse Correlated Equilibrium $\boldsymbol{s} \in \mathbb{BCCE}(\boldsymbol{V}, \mathrm{x}, \delta)$ *is a joint strategy profile that, for any bidder* $i \in [n]$*, value* $v \in \mathrm{supp}_i(\boldsymbol{v})$*, and deviation bid* $b^* \geq 0$*,*

$$\mathop{\mathbf{E}}_{\boldsymbol{v},\, \boldsymbol{s},\, \mathrm{x}}\big[\, u_i(v_i,\, \boldsymbol{s}(\boldsymbol{v})) \,\big|\, v_i = v \,\big] \;\geq\; \mathop{\mathbf{E}}_{\boldsymbol{v},\, \boldsymbol{s},\, \mathrm{x}}\big[\, u_i(v_i,\, b^*,\, \boldsymbol{s}_{-i}(\boldsymbol{v})) \,\big|\, v_i = v \,\big] - \delta.$$

It is well-known (see [Rou15]) that all equilibrium concepts together form the following inclusion: Bayesian Nash Equilibrium $\subseteq$ Bayesian Correlated Equilibrium $\subseteq$ Bayesian Coarse Correlated Equilibrium*.*

We remark that Definition A.2 follows the definition of Bayesian (Coarse) Correlated Equilibria by [CKK+15], different from those by [ST13] (see [Syr14, Chapter 3.3.1] for a thorough discussion).

THEOREM A.1. (TIGHT PoS) *Given a joint value distribution* $\boldsymbol{V} \in \mathbb{V}_{\mathsf{joint}}$ *and any tie-breaking rule* $\mathrm{x}^* \in \mathbb{FPA}$*, for any* $\delta > 0$*, there exists a joint strategy profile* $\boldsymbol{s}(\boldsymbol{v}) \equiv (s_i(\boldsymbol{v}))_{i \in [n]}$ *such that*

1. *The expected auction/optimal* Social Welfares *are equal* $\mathsf{FPA}(\boldsymbol{V}, \mathrm{x}^*, \boldsymbol{s}) = \mathsf{OPT}(\boldsymbol{V}, \mathrm{x}^*, \boldsymbol{s})$*.*

2. *It forms a* δ*-approximate* Bayesian Correlated Equilibrium $\boldsymbol{s} \in \mathbb{BCE}(\boldsymbol{V}, \mathrm{x}^*, \delta)$ *and thus also a* δ*-approximate* Bayesian Coarse Correlated Equilibrium $\boldsymbol{s} \in \mathbb{BCCE}(\boldsymbol{V}, \mathrm{x}^*, \delta)$*.*

Proof. Let us consider the first-order valuer $h = h(\boldsymbol{v}) := \mathrm{argmax}(\boldsymbol{v}) \in [n]$; breaking ties in favor of the smallest index. We explicitly construct a (deterministic) workable joint strategy profile:
(i) The first-order valuer h bids the second highest value, namely $s_h(\boldsymbol{v}) = \max(\boldsymbol{v}_{-h})$.
(ii) Each other bidder $i \in [n] \setminus \{h\}$ bids her value minus a δ term, namely $s_i(\boldsymbol{v}) = v_i - \delta$.
This strategy profile $\boldsymbol{s}$ always allocates the item to the first-order valuer h and realizes the optimal Social Welfare $= \max(\boldsymbol{v})$. In expectation, we have **Item 1** that $\mathsf{FPA}(\boldsymbol{V}, \mathrm{x}^*, \boldsymbol{s}) = \mathsf{OPT}(\boldsymbol{V}, \mathrm{x}^*, \boldsymbol{s})$.

The first-order valuer h realizes a utility $u_h(v_h, \boldsymbol{s}(\boldsymbol{v})) = v_h - s_h(\boldsymbol{v}) = \max(\boldsymbol{v}) - \max(\boldsymbol{v}_{-h}) \geq 0$. The threshold winning bid for this bidder h is the highest other bid $\max(\boldsymbol{s}_{-h}(\boldsymbol{v})) = \max(\boldsymbol{v}_{-h}) - \delta$. Thus with another deviation bid $b_h^* \geq 0$, bidder h realizes a deviation utility $u_h(v_h, b_h^*, \boldsymbol{s}_{-h}(\boldsymbol{v})) \leq (v_h - b_h^*) \cdot \mathbb{1}(b_h^* \geq \max(\boldsymbol{s}_{-h}(\boldsymbol{v}))) \leq \max(\boldsymbol{v}) - \max(\boldsymbol{v}_{-h}) + \delta = u_h(v_h, \boldsymbol{s}(\boldsymbol{v})) + \delta$, which is at most a δ increase over the current utility.

Each other bidder $i \in [n] \setminus \{h\}$ realizes a zero utility $u_i(v_i, \boldsymbol{s}(\boldsymbol{v})) = 0$. The threshold winning bid for this bidder i is the highest other bid $\max(\boldsymbol{s}_{-i}(\boldsymbol{v})) = s_h(\boldsymbol{v}) = \max(\boldsymbol{v}_{-h}) \geq v_i$. To win in the considered First Price Auction $\mathrm{x}^* \in \mathbb{FPA}$, bidder i must overbid $b_i^* \geq \max(\boldsymbol{s}_{-i}(\boldsymbol{v})) \geq v_i$ and realize a nonpositive deviation utility $u_h(v_h, b_h^*, \boldsymbol{s}_{-h}(\boldsymbol{v})) \leq 0$.

In sum, the considered strategy profile $\boldsymbol{s}$ forms a δ-approximate Bayesian Correlated Equilibrium $\boldsymbol{s} \in \mathbb{BCE}(\boldsymbol{V}, \mathrm{x}^*, \delta)$. **Item 2** follows then. This finishes the proof. ∎

Our δ-approximate Bayesian Correlated Equilibrium has almost the same output as the second-price auction (under the truthful strategies). We would like to present Theorem A.1 in terms of a universal δ-approximate Bayesian Correlated Equilibrium for any tie-breaking rule $\mathrm{x}^* \in \mathbb{FPA}$, instead of an exact equilibrium for a particular tie-breaking rule x that is compatible with the underlying value distribution $\boldsymbol{V}$. We would avoid the latter, because a typical tie-breaking rule should only depend on the *bid profile* and the bidders' identities, while a compatible tie-breaking rule further keeps track of the *value profile*.