Revenue Guarantee across Bayes Coarse Correlated Equilibria*

Linchen Zhang

July 22, 2024

Abstract

A single unit item owner sells the good to buyers through an auction design. We study the revenue guarantee (the lowest expected revenue amount that the seller can get) across all Bayes coarse correlated equilibria (BCCE, a weaker and inclusive equilibrium concept than Bayes correlated equilibrium, or BCE), starting with analysis in a first-price auction and a common value setting. We first show that it suffices to only examine identical-play equilibria, where all buyers follow exactly the same actions. Based on this simplification, we give characteristic functions for the revenue guarantee assuming a continuous value distribution. Revenue guarantee under BCCE converges to that under BCE when the market goes large. We also extend the model to allow for discrete distributions, other standard auctions, and symmetric prior environments.

*Linchen Zhang: Department of Economics, University of California San Diego, liz032@ucsd.edu.

1 Introduction

In mechanism designs and information economics, a key object of interest is the revenue, how much the seller would collect from sales in expectation. However, this number can vary to a huge extent depending on different information structures (that is, different signals received by the buyers) and on the equilibrium concept, as heterogeneous signals will generate dissimilar beliefs and different notions of equilibrium will give disparate equilibrium constraints on the buyers' actions. A natural question comes for the seller: given a certain equilibrium definition, what would be the minimum revenue she could get across all possible information structures? The answer to this question provides a worst-case analysis. This minimum revenue, defined as a revenue guarantee, guarantees the least amount of payoff the seller would obtain across all equilibria and is robust to information by construction. If the guarantee is high relative to the full potential payoff, then the seller would be more satisfied.

At first look, the revenue guarantee might be unsolvable unless we make additional assumptions on the exact forms of information structure, as most existing literature did, but that is not the case. Under the first-price auction (FPA), Bergemann, Brooks and Morris (2017) provides the characterization function of revenue guarantee across Bayes correlated equilibria (BCE) in a common-value single-item setting. Nonetheless, the concept of BCE might be a narrow definition: As Hartline, Syrgkanis and Tardos (2015) suggests, no-regret learning dynamic, a fundamental feature shared by many algorithms in applications, converges to Bayes coarse correlated equilibrium (BCCE), a broader and inclusive concept of the BCE. It would be useful to explore the revenue guarantee under the notion of BCCE, as it gives a more solid guarantee subject to a larger set of information structures and in practice may offer more insights on the least possible amount of revenue.

This paper analyzes the revenue guarantee problem under the notion of BCCE. We start our analysis in the first-price auction with a common-value item, and then generalize the results from FPA to any standard auctions with no losing payments and from common values to symmetric priors. In the paper, we first show that to study the revenue guarantee problem under BCCE, it suffices to examine only identical play BCCEs, that is, when all the buyers act in the same way, and by that, it is without loss of generality to consider them jointly as a single representative buyer. This change allows us to disregard nuances between buyers' action profiles and thus largely simplifies the problem. Then based on the simplification, we characterize revenue guarantee when it takes a continuous distribution and when the value takes a binary distribution of 0 and 1. The asymptotic revenue guarantee of a continuous distribution is equal to the revenue guarantee under BCE characterized by Bergemann, Brooks and Morris (2017). The result of binary distribution also offers a lower bound on the revenue guarantee of any other arbitrary distribution of the same expected value.

Our research relates to literature in four strands: the first-price auction, revenue guarantee, Bayes coarse correlated equilibrium, and no-regret learning dynamics. Firstly, our work is related to a large area of literature on the FPA: from earlier papers that make additional assumptions on information structures such as Vickrey (1961), Milgrom and Weber (1982), Battigalli and Siniscalchi (2003), Dekel and Wolinsky (2003), Fang and Morris (2006), Azacis and Vida (2015), and Feldman, Lucier and Nisan (2016) to the recent work by Bergemann, Brooks and Morris (2017) that allows for information variations but examines under the notion of BCE. Our paper advances the analysis and characterizes the revenue guarantee under a weaker notion of BCCE. It is also worth noting that this paper's results apply beyond FPA: it holds for any standard auctions without losing payments. Secondly, it contributes to the literature on BCCE. Proposed by Moulin and Vial (1978), though an important concept to which many learning algorithms converge, BCCE has been barely studied in the economics literature. It is mainly because there are no simple optimal conditions like the first-order condition that one can start analyzing with. Also, the equilibrium concept is so coarse that very few affirmative arguments concerning the equilibrium outcomes can be declared. However, as we show in this paper, when combining BCCE with the worst-case scenario analysis, it may be possible to derive certain results on the equilibrium outcomes. Thirdly, this paper contributes to the literature on revenue guarantee such as Bergemann, Brooks and Morris (2017), Du (2018), and Brooks and Du (2021) in that it is one of the first papers that investigates revenue guarantee beyond the scope of BCE. There are also papers studying the guarantee in terms of welfare analysis, namely the price of anarchy, such as Syrgkanis and Tardos (2013) and Jin and Lu (2023), but that is out of the frame of this paper. Lastly, this paper resonances with no-regret learning literature such as Hart and Mas-Colell (2000) and Blum et al. (2008), and as Hartline, Syrgkanis and Tardos (2015) suggests, no-regret learning dynamics always converge to a solution concept of BCCE. Therefore, our analysis under BCCE may offer insights into those industrially applied no-regret algorithmic design outcomes.

The rest of the paper proceeds as follows. Section 2 introduces the model framework of the revenue guarantee problem. Section 3 shows that it suffices to study only with identical play BCCEs, arguing first with 2 buyers and then generalizing to N buyers. Section 4 characterizes the revenue guarantee for continuous distributions. Section 5 analyzes a binary distribution revenue guarantee, which also serves as a lower bound of the revenue guarantee for all distributions with the same expected value. Section 6 extends the analysis to standard auctions and symmetric priors. Section 7 concludes and discusses potential applications.

2 Model

A principal who owns a single-unit good intends to sell the item to buyers indexed by $i \in \{1, ..., N\}$ through the first-price auction. The good has a common value $v \in V = [\underline{v}, \overline{v}] \subseteq [0, \infty)$, distributed according to a probability measure $\mu \in \Delta(V)$. In this section (and up to section 6), for simplicity, we restrict the mechanism of sale to be the first-price auction and the item to be commonly valued. We will later show in section 6, these two restrictions can be easily extended to standard auctions and symmetric prior.

The action profile is denoted by $A = \prod_{i=1}^{N} A_i \subseteq V^N$. The buyers' private information about the state is described by an information structure $I = (S, \omega)$, where S_i is a set of signals for buyer i, $S = \prod_{i=1}^{N} S_i$, and $\omega \in \Delta(V \times S)$ is the joint distribution of the states and signals such that $\operatorname{marg}_V \omega = \mu$.

Define a strategy for buyer i to be a mapping from signals to actions $\alpha_i : S_i \to \Delta(A_i)$, then

$$da = \int_{s \in S} \prod_i \alpha_i(a_i|s_i) ds.$$

Define $\tilde{\omega} \in \Delta(V \times S \times A)$ to be the joint distribution of values, signals, and actions, where the actions are constructed based on the strategy α and the info structure ω. In other words, $\omega = \operatorname{marg}_{\{V,S\}} \tilde{\omega}$.

Then let us restrict our attention to only the values and actions: Define $\sigma \in \Delta(V \times A)$ to be the marginal joint distribution of values and signals, or $\sigma = \text{marg}_{\{V,A\}} \tilde{\omega}$. From now on, we will work directly in the σ world. It offers us a convenience to ignore all the signal structures, as signals do not directly affect the equilibrium constraints under BCCE. Nonetheless, signals are still implicitly used as a way for buyers to correlate their actions.

Therefore, buyer i's expected utility given σ is

$$U_i(\sigma) = \int_{v \in V} \int_{a \in A} u_i(v, a) \sigma(dv, da).$$

In the first-price auction setting, let $X = \max_j a_j$ and $|J|$ be the number of buyers who bid the highest price, then the utility function of buyer i given a specific realization of v and a concrete action profile a is

$$u_i(v, a_i, a_{-i}) = \mathbb{1}_{a_i = M} \times \frac{1}{|J|} \times (v - X).$$

In other words, if there is only one buyer who bids for the highest value, then she will win the item and pay the price she offers. If there are more than 1 buyers who offer the highest bid, they will get a share of the item and pay the corresponding fractional price if the item is divisible, or get the entire item and pay full price with $1/|J|$ probability. In the latter case, we consider all buyers to be risk-neutral.

A Bayes correlated equilibrium (BCE) is a joint distribution σ such that:

$\forall i,\ \forall a_i',$

$$\int_{v \in V} \int_{a_{-i} \in A_{-i}} u_i(v, a_i, a_{-i}) \sigma(dv, da) \geq \int_{v \in V} \int_{a_{-i} \in A_{-i}} u_i(v, a_i', a_{-i}) \sigma(dv, da).$$

In another sense, BCE needs to make sure that, conditional on signals s_i that is associated with a_i, any deviation a_i' is not preferred, so it optimizes conditional on a_i.

On the other hand, a Bayes coarse correlated equilibrium (BCCE) is a joint distribution σ such that: $\forall i,\ \forall a_i',$

$$\int_{v \in V} \int_{a \in A} u_i(v, a_i, a_{-i}) \sigma(dv, da) \geq \int_{v \in V} \int_{a \in A} u_i(v, a_i', a_{-i}) \sigma(dv, da).$$

In contrast, BCCE only needs to satisfy that any fixed deviation a_i' that doesn't depend on s_i is unfavorable. Thus, it optimizes ex-ante (not conditional on a_i).

Let us consider an example to illustrate the difference of those two concepts. Suppose it is still in the first-price auction, and the item value v is uniformly distributed in $[0, 1]$, and all buyers i receive a perfect information $s_i = v$. Then consider the strategy $a_i = 0.5s_i$, so all buyers bid half of the item value. It is not a BCE, as some buyer i could bid a bit more than half of the value and win the item entirely. However, it is a BCCE: suppose a buyer i deviates to a fixed action δ, she would win the item with values ranged in $[0, 2\delta]$ and pay δ, so the expected payoff is 0, which is less than the positive payoff following equilibrium.

It is also straightforward to observe that BCCE is a weaker concept than BCE. For buyer i, BCE constraints must hold for all distribution conditional on a_i, and integrating across a_i directly implies the BCCE constraint for buyer i. This is where the name "coarse" comes from.

Given the common value restriction, define

$$\mathbb{E}(v) = \int_v v\mu(dv).$$

Then the expected revenue can be expressed as

$$\begin{aligned} R(\sigma) &= \mathbb{E}(v) - \sum_i \mathbb{E}(u_i) \\ &= \mathbb{E}(v) - \sum_i \int_{v \in V} \int_{a \in A} u_i(v, a)\sigma(dv, da). \end{aligned}$$

We define the revenue guarantee of a mechanism M as the minimum expected revenue over all BCCE σ:

$$RG = \min_\sigma R(\sigma), \text{ where } \sigma \text{ is a BCCE.}$$

We call a BCCE σ that attains the revenue guarantee level RG as a min-R BCCE. The purpose of this paper is to find such RG. We will later construct an example that achieves the minimum revenue level, so the minimum is well-defined.

Finally, let us remind ourselves that after solving the problem over σ, we need to construct an information structure and corresponding bidding strategies to sustain those action profiles.

3 Identical Plays

In this section, we show that to solve the revenue guarantee problem under FPA, it suffices to examine only identical-play equilibria, where all buyers' actions are exactly the same. In another sense, we can view all buyers as a single representative buyer who jointly bids according to a single strategy. This result would largely simplify the analysis, as it is no longer needed to investigate the interactions and nuances between buyers. Note that "identical plays" is a stronger assertion than "symmetric strategies", as the later one may yield different joint actions given heterogeneous signals.

3.1 Two Buyers

We begin the arguments with two buyers and then generalize the results to N buyers. The case of two buyers is more intuitive, and it incorporates the key idea of the analysis.

For the purpose of being manipulated later, information of μ and σ is stored in two dictionaries, where v and (v, a) are the keys respectively, and the corresponding probabilities can be accessed by calling $\mu(v)$ and $\sigma(v, a)$. For two buyers, $(v, a) = (v, a_1, a_2)$. The continuous distribution of μ and σ can be well approximated by an infinitesimal grid size of v and a, and then it is possible to store that finite information in the dictionary. This design allows us to easily modify probability values for the keys in the info structure σ. For instance, if we want to replace (v, c_1, c_2) with a probability p to (v, d_1, d_2) with probability $p/2$ and (v, e_1, e_2) with probability $p/2$, we just need to subtract p for the original probability value stored in (v, c_1, c_2), and add $p/2$ to both values stored in (v, d_1, d_2) and (v, e_1, e_2). However, recall that for σ to be a valid information structure, it has to be that $\sum_a \sigma(v, a) = \mu(v)$ for all possible v and that $\sum_{v,a} \sigma(v, a) = 1$. These are the two conditions we need to check after modifications on σ. For the example above, suppose $\tilde{\sigma}$ is the joint distribution after changes and σ is the one before changes, we write the modification as

$$\tilde{\sigma} = \sigma - p \times (v, c_1, c_2) + \frac{p}{2} \times (v, d_1, d_2) + \frac{p}{2} \times (v, e_1, e_2).$$

This construction method may be seen several times throughout proofs.

Proposition 1. *With only 2 buyers, the min-R BCCE in FPA has only identical plays (i.e., if σ is a min-R BCCE, and for any (v, a_1, a_2) such that $\sigma(v, a_1, a_2) > 0$, then it must be that $a_1 = a_2$.)*

Proof. Let $\tilde{\sigma}$ be a mirror distribution of σ with a switch of index. By this construction, buyer 1 in $\tilde{\sigma}$ behaves exactly the same as buyer 2 in σ, and buyer 2 in $\tilde{\sigma}$ the same as buyer 1 in σ. For any tuple (v, a_1, a_2) such that $\sigma(v, a_1, a_2) > 0$, suppose $\sigma(v, a_1, a_2) = p$, then $\tilde{\sigma}(v, a_2, a_1) = \sigma(v, a_1, a_2) = p$. Since σ is a BCCE, then by symmetry, $\tilde{\sigma}$ is also a BCCE and $R(\sigma) = R(\tilde{\sigma})$.

Now consider $\widehat{\sigma} = \frac{1}{2}\sigma + \frac{1}{2}\tilde{\sigma}$ (with 1/2 probability it follows σ, and with 1/2 probability it follows $\tilde{\sigma}$; $\widehat{\sigma}$ can be easily constructed using the dictionaries).

By the linearity of expectation, $R(\widehat{\sigma}) = \frac{1}{2}R(\sigma) + \frac{1}{2}R(\tilde{\sigma}) = R(\sigma)$. By the convexity of BCCE, $\widehat{\sigma}$ is still a BCCE. The convexity of BCCE is directly implied as the BCCE equilibrium constraint is linear. Combining these two we get that $\widehat{\sigma}$ is also a minimum revenue BCCE. Now let

$$\widehat{\sigma}(v) = \frac{1}{2}\sigma(v, a_1, a_2) \times (v, a_1, a_2) + \frac{1}{2}\tilde{\sigma}(v, a_2, a_1) \times (v, a_2, a_1) = \frac{p}{2} \times (v, a_1, a_2) + \frac{p}{2} \times (v, a_2, a_1).$$

Consider a slightly altered distribution

$$\gamma = \widehat{\sigma} - \widehat{\sigma}(v) + \frac{p}{2} \times (v, a_1, a_1) + \frac{p}{2} \times (v, a_2, a_2).$$

In other words, we modify the dictionary $\widehat{\sigma}$ and construct a new dictionary γ in the following way: first delete the $\widehat{\sigma}(v)$ parts of $\widehat{\sigma}$ (subtract probability value $p/2$ in the key of (v, a_1, a_2) and probability $p/2$ in the key of (v, a_2, a_1)), then add probability value $p/2$ each to keys (v, a_1, a_1) and (v, a_2, a_2). It is easy to check conditions that $\sum_a \sigma(v, a) = \mu(v)$ for all possible v and that $\sum_{v,a} \sigma(v, a) = 1$ are satisfied, so this modification is valid.

Now let us compare γ with $\widehat{\sigma}$. First we show that $R(\widehat{\sigma}) \geq R(\gamma)$. By the nature of the FPA pricing function,

$$R(\widehat{\sigma}) - R(\gamma) = p \times \max(a_1, a_2) - p/2 \times (a_1 + a_2) \geq 0.$$

Consequently, since buyers' strategies are both symmetric in $\widehat{\sigma}$ and γ, then $\forall i$,

$$E_{\widehat{\sigma}}(u_i) = \frac{\mathbb{E}(v) - R(\widehat{\sigma})}{2} \leq \frac{\mathbb{E}(v) - R(\widehat{\gamma})}{2} = \mathbb{E}_{\gamma}(u_i)$$

Secondly, we show that γ is also a BCCE. Let $\widehat{\sigma}_{rest} = \widehat{\sigma} - \widehat{\sigma}(v)$, the common part shared in $\widehat{\sigma}$ and γ. Then from the perspective of BCCE equilibrium constraint, since $\widehat{\sigma}$ is a valid BCCE (we use a different notation b as a is fixed in the previous argument),

$$\forall i,\ \forall b_i',\int_{v,b_i,b_{-i}} u_i(v,b_i,b_{-i})d\widehat{\sigma}(v,b_i,b_{-i}) \geq \int_{v,b_i,b_{-i}} u_i(v,b_i',b_{-i})d\widehat{\sigma}(v,b_i,b_{-i}).$$

Here we replace the $\sum$ with $\int$ as we can always take the grid size of the dictionary to be arbitrarily small. Rewrite the equation,

$$\begin{aligned}\forall i,\ \forall b_i',\ \mathbb{E}_{\widehat{\sigma}}(u_i) &\geq \int_{v,b_i,b_{-i}} u_i(v,b_i',b_{-i})d\widehat{\sigma}(v,b_i,b_{-i}) \\ &= \int_{v,b_i,b_{-i}} u_i(v,b_i',b_{-i})d\widehat{\sigma}_{rest}(v,b_i,b_{-i}) + \frac{p}{2}\times u_i(v,b_i',a_1) + \frac{p}{2}\times u_i(v,b_i',a_2) \\ &= \int_{v,b_i,b_{-i}} u_i(v,b_i',b_{-i})d\gamma(v,b_i,b_{-i}).\end{aligned}$$

Therefore $\forall i,\ \forall b_i'$,

$$\begin{aligned}\mathbb{E}_{\gamma}(u_i) &\geq \mathbb{E}_{\widehat{\sigma}}(u_i) \\ &= \int_{v,b_i,b_{-i}} u_i(v,b_i',b_{-i})d\gamma(v,b_i,b_{-i}).\end{aligned}$$

Hence, γ is also a valid BCCE. However, because $\widehat{\sigma}$ is a minimum revenue BCCE, it must be that $R(\widehat{\sigma}) \leq R(\gamma)$. This is only allowed to occur when $\max(a_1,a_2) = \frac{1}{2}(a_1+a_2)$, or when $a_1 = a_2$, then it must be that $a_1 = a_2$. □

The economic intuition is that identical bids can reduce price competition among buyers, and that perfect collusion leads to a highest total payoff for buyers and a lowest revenue for the seller.

3.2 N Buyers

An interesting fact in the setting of 2 buyers is that the plays must be identical for a minimum revenue BCCE. However, this result does not need to hold when $N > 2$. For more than two buyers in the FPA, the only guarantee is that the first max must be the same as the second max, or the highest bid must be the same as the second highest bid, and it has no restriction on the rest of the bids. Nonetheless, it is without loss of generality to consider all the rest bidding the same as the first max buyer.

If we go beyond the first-price auction to the standard auction setting, then we need to replace the "second max" to a "marginal type" suitable in the auction setting, so that the "first max" (or the winning type) is the same as the "marginal type" in the minimum revenue case. We need to assume positive payoffs for all buyer to eliminate uninteresting cases, and we will talk more about it in section 6.

Proposition 2. *Let the first max (the highest number in a set) be* $\max_1$*, the second max (the second highest number in a set) be* $\max_2$*, then under FPA, if* σ *is a min-R BCCE, for any* $(v, a_1, a_2, ..., a_n)$ *such that* $\sigma(v, a_1, a_2, ...a_n) > 0$*, it must be that* $\max_1\{a_1, ..., a_n\} = \max_2\{a_1, ..., a_n\}$

Proof. See Appendix. □

Lemma 1. *(Identical plays) There exists a min-R BCCE* σ *such that for all* $\sigma(v, a_1, a_2, ..., a_n) > 0$*,* $\exists \alpha$ *s.t.* $a_1 = a_2 = ... = a_n = \alpha$*. We call such* σ *an identical play BCCE.*

Proof. See Appendix. □

4 Characterization of a Continuous Distribution

This section characterizes the revenue guarantee problem where $\mu(v)$ is required to be an absolutely continuous distribution, so that a probability density function exists. For simplicity, the range of v is normalized to $[0, 1]$. We list our main characterization results first, and then dive into all the technical details. For derivations, We start with the asymptotic problem, as it is more intuitive to approach and has a closed-form solution, and then attack the problem of a finite number of buyers, where the revenue guarantee is characterized by a set of equations.

Theorem 1. *The revenue guarantee level RG across BCCE (*RG_{BCCE}*) with an associated threshold value* $\widehat{v}$ *is characterized by the following equations,*

$$\begin{cases} \frac{\mathbb{E}(v) - RG_{BCCE}}{N} = \int_{v=0}^{\widehat{v}} v f(v) dv, \\ RG_{BCCE} = \int_{v=0}^{1} H(v) f(v) dv, \end{cases} \tag{1}$$

where

$$H(v) := \mathbb{E}(v'|\widehat{v} \leq v' \leq v) \times \frac{F(v) - F(\widehat{v})}{F(v)}, \tag{2}$$

and F and f are the cdf and pdf of μ, $\mu(dv) = f(v)dv$.

Moreover, $\lim_{N\to\infty} RG_{BCCE} = \lim_{N\to\infty} RG_{BCE}$. In a large market, the revenue guarantee across BCCE and that across BCE are the same.

Based on this theorem, we could solve for the RG_{BCCE} and $\widehat{v}$. Let us now discuss why a unique solution exists. We first show the existence of a solution. Let us reduce to one equation with only $\widehat{v}$ from the system of equations in thorem 1:

$$\mathbb{E}(v) = \int_{v=0}^{1} \left[\int_{m=\widehat{v}}^{v} mf(m)dm\right] \frac{f(v)}{F(v)}dv + N\int_{v=0}^{\widehat{v}} vf(v)dv$$

When $\widehat{v} = 0$, the right-hand side $\leq \mathbb{E}(v)$. When $\widehat{v} = 1$, the right-hand side $\geq N\mathbb{E}(v) > \mathbb{E}(v)$. Then by Intermediate Value Theorem, there exists a $\widehat{v} \in [0, 1]$ that satisfies the right-hand side $= \mathbb{E}(v)$, so a solution always exists. Furthermore, as $\widehat{v}$ increases, RG_{BCCE} decreases, so the solution must be unique by construction, otherwise RG_{BCCE} won't be the minimum revenue level.

4.1 Simplified Setup

Before diving into computational details, let us first simplify the problem using what we have discovered so far. Given the identical play argument, we can think of all buyers jointly as a single agent, and then the revenue minimization problem can be simplified as (here a is one-dimensional)

$$\min_{\sigma} \int a \times \sigma(dv, da),$$

which is subject to the equilibrium constraint that for any possible deviation δ,

$$\frac{\mathbb{E}(v) - RG_{BCCE}}{N} \geq \int \mathbb{1}(\delta > a)(v - \delta)\sigma(dv, da), \tag{3}$$

and that σ is an identical play min-R BCCE that satisfies $\int_a \sigma(v, da) = \mu(v)$ for all v. The left-hand side is the utility obtained in the equilibrium, and the right-hand side is the utility amount when deviating to δ for buyer i, as he would win the item when $\delta > a$ and pay the price of δ.

This obedience constraint restricts that any deviation δ is not favored to following the equilibrium strategies.

To further describe the relation between a and v in σ, we introduce the definition of monotonicity, so that as v increases, a also increases, and we may write one as a function of the other.

Lemma 2. *(Monotonicity) Without loss of generality, we can treat the identical play min-R BCCE σ as monotonic, that is $\forall(v_1, a_1), (v_2, a_2)$ such that $\sigma(v_1, a_1) > 0, \sigma(v_2, a_2) > 0$, if $v_1 > v_2$, then $a_1 > a_2$.*

Proof. See Appendix. □

The intuition is that for the minimum revenue scenarios, it is never optimal to pay a higher price on an item with less value. Based on this lemma, we can treat σ as a monotonic identical play min-R BCCE. Then we can define the function

$$\overline{v}(\delta) = \sup\{v : \sigma(v, a) > 0, \delta > a\}$$

that finds the highest possible value v where a potential deviation play δ could win the item. The main purpose of introducing $\overline{v}$ is to reduce the 2-D σ space of (v, a) to 1-D μ space of v, which is known and can be numerically tracked.

4.2 Asymptotic Characterization

Let us look at a special case, when the number of buyers is arbitrarily large. This asymptotic case embeds essential proof ideas for the general case in theorem 1. To start, we copy the equilibrium constraint from equation 3 here:

$$\frac{\mathbb{E}(v) - RG_{BCCE}}{N} \geq \int \mathbb{1}(\delta > a)(v - \delta)\sigma(dv, da),$$

As $N \to \infty$, the limit of the left-hand side

$$\lim_{N\to\infty} \frac{\mathbb{E}(v) - RG_{BCCE}}{N} = 0,$$

and the right-hand side does not depend on N. So for the asymptotic case, when the number of buyers goes large, the equilibrium constraint can be further simplified as $\forall\delta$,

$$0 \geq \int \mathbb{1}(\delta > a)(v-\delta)\sigma(dv, da).$$

Moreover, using $\overline{v}(\delta)$, the equilibrium constraint can be written as for all potential deviation δ,

$$0 \geq \int_{v=0}^{\overline{v}(\delta)} (v-\delta)\mu(dv).$$

Finally, observe that when $\delta = 0$, the constraint is binding. If any constraint of $\delta \in (0, \mathbb{E}(v)]$ is slack, then we can decrease bidding values a associated with $\overline{v}(\delta)$ and at the same time force the equilibrium inequality condition to reach an equality level. As this change still satisfies all the constraints, it is still a BCCE, but now it achieves a lower revenue amount. The upper bound $\mathbb{E}(v)$ is chosen since it is never profitable to deviate to a level $\delta > \mathbb{E}(v)$ where the potential deviation would win the item in all cases and get a surely negative payoff. Hence, for a minimum revenue BCCE, all the constraints must be binding for $\delta \in [0, \mathbb{E}(v)]$. Fix a δ, we have

$$0 = \int_{v=0}^{\overline{v}(\delta)} (v-\delta)\mu(dv),$$

or equivalently,

$$\delta F(\overline{v}(\delta)) = \int_{v=0}^{\overline{v}(\delta)} v f(v) dv.$$

Hence, it must be that $\forall\delta \in [0, 1]$,

$$\delta = \mathbb{E}[v | v \leq \overline{v}(\delta)].$$

Let function $G(v) := \mathbb{E}[v' | v' \leq v]$. Now this offers another angle why $\delta \in [0, \mathbb{E}(v)]$, as the conditional mean $G(v)$ cannot exceed $\mathbb{E}(v)$. Since μ is known and continuous, $G(v)$ can be calibrated. G^{-1} also exists and can be computed as G is strict increasing in v. Writing in terms of G^{-1}, $\overline{v}(\delta) = G^{-1}(\delta)$ is also characterized. Finally, note that $\overline{v}$ reveals the relation between v and a, so $v(a) = \overline{v}(a) =$

$G^{-1}(a)$ and $a = G(v)$. Therefore, the corresponding revenue guarantee (across all BCCEs) is

$$\begin{aligned} RG_{BCCE} &= \int_{a=0}^{\mathbb{E}(v)} \int_{v=0}^{1} a\sigma(dv, da) \\ &= \int_{v=0}^{1} G(v)\mu(dv) \\ &= \int_{v=0}^{1} G(v)f(v)dv. \end{aligned}$$

This characterizes the asymptotic revenue guarantee given a continuous distribution μ.

We can compare our result to the existing literature (Bergemann, Brooks and Morris (2017)), which has an asymptotic revenue guarantee across all BCEs

$$RG_{BCE} = \int_{v=0}^{1} -v\ln(F(v))f(v)dv.$$

It turns out the revenue guarantee under BCE and the revenue guarantee under BCCE are equal asymptotically. To see the equivalence result, let $x = \int_{t=0}^{v} tf(t)dt$, $y = \ln(F(v))$, then

$$RG_{BCCE} - RG_{BCE} = \int xdy + \int ydx = \Big[xy\Big]\Big|_{v=0}^{1} = 0.$$

As we enlarge our definition of equilibria, we happen to not lose any potential worst-case revenues asymptotically. However, as we will show later, a difference in the revenue guarantees exists for finite N, but the gap shrinks to zero as N goes large.

To get a better sense, let us look at an example here. Based on the characteristic function above, we can calculate the asymptotic revenue guarantee of FPA across BCCE under a uniform distribution. Given $\mu(v) \sim Uniform[0, 1]$, we know

$$G(v) = \mathbb{E}[v'|v' \leq v] = \frac{v}{2}, \text{ and } f(v) = 1,$$

so

$$RG_{BCCE} = RG_{BCE} = \int_{v=0}^{1} \frac{1}{2}vdv = \frac{1}{4}.$$

4.3 Finite N Characterization

Now we look beyond the asymptotic problem and focus on a finite number of buyers. The only difference between a finite and infinite N is that now the equilibrium constraint is more relaxed, and a potential deviation δ is preferred only when it has more payoff than the amount of payoff shared between buyers. Mathematically, for any possible deviation δ,

$$\frac{\mathbb{E}(v) - R}{N} \geq \int \mathbb{1}(\delta > a)(v - \delta)\sigma(dv, da),$$

or in $\overline{v}$ measure,

$$\frac{\mathbb{E}(v) - R}{N} \geq \int_{v=0}^{\overline{v}(\delta)} (v - \delta)\mu(dv).$$

Still using the same approach as the asymptotic case, to lower revenue, we want the constraint to be binding as much as possible. Therefore, for action a associated with small item values v, we assign $a = 0$ to them, up to a cutoff value $\widehat{v}$ that satisfies

$$\frac{\mathbb{E}(v) - R}{N} = \int_{v=0}^{\widehat{v}} v\mu(dv). \tag{4}$$

For $v > \widehat{v}$, as the equilibrium constraint must be binding,

$$\frac{\mathbb{E}(v) - R}{N} = \int_{v=0}^{\overline{v}(a)} (v - a)\mu(dv),$$

so

$$aF(\overline{v}(a)) + \int_{v=0}^{\widehat{v}} vf(v)dv = \int_{v=0}^{\overline{v}(\delta)} vf(v)dv,$$

then

$$a = \mathbb{E}(v|\widehat{v} \leq v \leq \overline{v}) \times \frac{F(\overline{v}(a)) - F(\widehat{v})}{F(\overline{v}(a))}.$$

Hence, let

$$H(v) := \mathbb{E}(v'|\widehat{v} \leq v' \leq v) \times \frac{F(v) - F(\widehat{v})}{F(v)},$$

then $a = H(v)$, and the corresponding revenue guarantee

$$RG_{BCCE} = \int_{v,a} a\sigma(dv, da) = \int_{v=0}^{1} H(v)f(v)dv \tag{5}$$

Now we have two equations ((4) and (5)) and two unknowns (R and $\widehat{v}$), so we could solve for both, and R is the revenue guarantee we are interested in. Note here the asymptotic problem is a special case of the general N problem: it has a restriction of $\widehat{v} = 0$ so that $H(v) = G(v)$. Finally, let us recall that we need to construct the signal structures that support the level of revenue guarantee. It turns out not to be a hard task: let $s_i = v$ and $a = H(s_i)$ for all i, and it attains the revenue guarantee level. This concludes our proof of Theorem 1.

Let us still study the uniform distribution example and consider $N = 2$. Now equation (4) gives

$$\int_{v=0}^{\widehat{v}} vdv = \frac{\frac{1}{2} - R}{2}, \text{ so } \widehat{v} = \sqrt{\frac{1}{2} - R}.$$

Combine it with equation (5),

$$\begin{aligned} R &= \int_{v=0}^{1} \frac{(v + \overline{v})}{2} \frac{(v - \overline{v})}{v} dv \\ &= \left[\frac{1}{4}v^2 - \frac{1}{2}\widehat{v}^2 \ln(v)\right]_{v=\widehat{v}}^{1} \\ &= \frac{1}{4} - \frac{1}{4}(\frac{1}{2} - R) + \frac{1}{2}(\frac{1}{2} - R) \ln \sqrt{\frac{1}{2} - R}. \end{aligned}$$

Solving it we get $R = 0.04708$. In contrast with a BCE revenue guarantee of $1/6$ from Bergemann, Brooks and Morris (2017), the BCCE revenue guarantee is at a much lower level. In fact, as N increases, the gap between the BCE and BCCE revenue guarantees reduces and eventually vanishes as $N \to \infty$.

5 Unknown Distributions

So far, we have characterized the BCCE revenue guarantee for a known continuous distribution μ. However, sometimes the seller may not know the exact distribution of μ, except for an expected item value $\mathbb{E}(v)$. In such cases, she may still be interested in the worst scenario payoff across all information structures and across all distributions, as it gives some ideas of the payoff level. We define this slightly different problem as the revenue guarantee problem given a fixed expected value. This section solves this extended problem. We first show that to find the revenue guarantee across all BCCEs under a fixed $\mathbb{E}(v)$, it is sufficient to only look at the binary distribution. We then attack the binary distribution problem using the same technical approach. Naturally, this offers a lower bound of revenue guarantee problem for any other known distribution μ with $E_\mu(v) = \mathbb{E}(v)$, and it doesn't require μ to be continuous. Therefore, it also offers some insights into the revenue guarantee when our continuous characterization method fails.

5.1 Binary Restrictions and Lower Bounds

Proposition 3. *(Lower bound for binary restriction) Under FPA,*

$$\min_{\mu(v)\ \ s.t. E_\mu(v)=\mathbb{E}(v)} \min_\gamma R(\gamma)|\mu \Longleftrightarrow \min_\gamma R(\gamma)|\underline{\mu}$$

where γ is a BCCE that satisfies $\sum_a \gamma(v,a) = \mu(v)$ for all v, and $\underline{\mu}$ is a binary distribution of v where $\underline{\mu}(0) = 1 - \mathbb{E}(v)$, and $\underline{\mu}(1) = \mathbb{E}(v)$. Note that here $\mathbb{E}_{\underline{\mu}}(v) = \mathbb{E}(v)$.

In other words, if Proposition 3 holds, it is sufficient to only restrict our attention to binary distribution $\underline{\mu}$ when solving the revenue minimization problem with any distribution μ such that the expected value is equal to $\mathbb{E}(v)$. Furthermore, for any value distribution μ with the same expected value, results from the binary restriction $\underline{\mu}$ give a lower bound for the revenue guarantee under μ.

Proof. See Appendix. □

5.2 Binary Asymptotic Results

By the identical play simplification, the revenue minimization problem can be written as

$$\min_\sigma \int a \times d\sigma(v,a),$$

which is subject to the equilibrium constraint that for any possible deviation δ,

$$\frac{1}{N}\int_{v,a}(v-a)\sigma(dv,da) \geq \int_{\delta>a}(v-\delta)\sigma(dv,da)$$

and that σ is an identical play BCCE consist with μ.

When $N \to \infty$, the revenue minimization problem is subject to

$$0 \geq \int_{\delta > a} (v - \delta) \sigma(dv, da).$$

Proposition 4. *If σ is a min-R BCCE such that $\forall (v, a_1, a_2, ..., a_n)$ where $\sigma(v, a_1, a_2, ..., a_n) > 0$, $a_1 = a_2 = ... = a_n = \alpha$ for some α, then $\alpha \leq v$. Moreover, there exists a min-R BCCE such that $\forall (v, a_1, a_2, ..., a_n)$ where $\sigma(v, a_1, a_2, ..., a_n) > 0$, $a_1 = a_2 = ... = a_n = \alpha \leq v$ for some α.*

Proof. See Appendix. ☐

For the binary distribution $\underline{\mu}$ with expected value $\mathbb{E}(v)$, by Proposition 4, for value $v = 0$, it is without loss to claim that $a = 0$, so the constraint can be written as

$$0 \geq (1 - \mathbb{E}(v))(0 - \delta) + \int \mathbb{1}(\delta > a)(1 - \delta) d\sigma(1, a).$$

Let $F(a)$ be the cdf of a given $v = 1$, then

$$F(\delta) = Prob(\delta > a | v = 1) = \frac{\int \mathbb{1}(\delta > a) d\sigma(1, a)}{\mathbb{E}(v)},$$

so the equilibrium constraint becomes

$$F(\delta) \leq \frac{\delta}{(1 - \delta)} \frac{(1 - \mathbb{E}(v))}{\mathbb{E}(v)}.$$

Observe that $F(\delta)$ is strictly increasing in δ, and as F is a cdf, $F(\delta) \leq 1$, so the constraint is effective only when

$$\delta \leq \mathbb{E}(v).$$

To minimize revenue, we want to lower the value of a as much as possible, or the highest F as much as possible, and that is achieved when every constraint is binding. Also note that when $\delta = 0$, the constraint is satisfied, and equality is reached. Therefore, the distribution of a given $v = 1$ is

$$f(a) = F'(a) = \frac{(1 - \mathbb{E}(v))}{\mathbb{E}(v)(1 - a)^2}, \text{ where } a \leq \mathbb{E}(v),$$

and

$$R = \int a d\sigma(1, a) = \int_0^{\mathbb{E}(v)} (1 - \mathbb{E}(v)) \frac{a}{(1 - a)^2} da = \mathbb{E}(v) + (1 - \mathbb{E}(v)) \ln(1 - \mathbb{E}(v)).$$

This solves the revenue guarantee across all BCCEs given $\mathbb{E}(v)$ and draws a lower bound on revenue across all distribution μ with an expected value of $\mathbb{E}(v)$. Note that here the full potential revenue is $\mathbb{E}(v)$, so the revenue loss (the part has to be given to all buyers) is

$$R_{loss} = \mathbb{E}(v) - R = -(1 - \mathbb{E}(v)) \ln(1 - \mathbb{E}(v)),$$

and the share of item value distributed to the seller is

$$Share_{Seller} = \frac{R}{\mathbb{E}(v)} = 1 + \frac{(1 - \mathbb{E}(v)) \ln(1 - \mathbb{E}(v))}{\mathbb{E}(v)}.$$

When $\mathbb{E}(v) = 1/2$, this gives us a revenue guarantee of $(1 - \ln 2)/2 = 0.1534$. If we compare it with the asymptotic revenue guarantee of the uniform distribution (which also has an $\mathbb{E}(v) = 1/2$), we will find it is much lower than the $1/4$ guarantee level of the uniform distribution. This finding is consistent with our lower-bound arguments.

In addition, if we plot the function of the seller's share, we can observe that it has a value of 0 when $\mathbb{E}(v) = 0$, strictly increases as $\mathbb{E}(v)$ rises, and eventually obtains a full share 1 when $\mathbb{E}(v)$ reaches 1. In other words, with a more right-skewed shape of the value distribution, there are more advantages on the seller's side, and she would obtain more shares of the total surplus against all buyers.

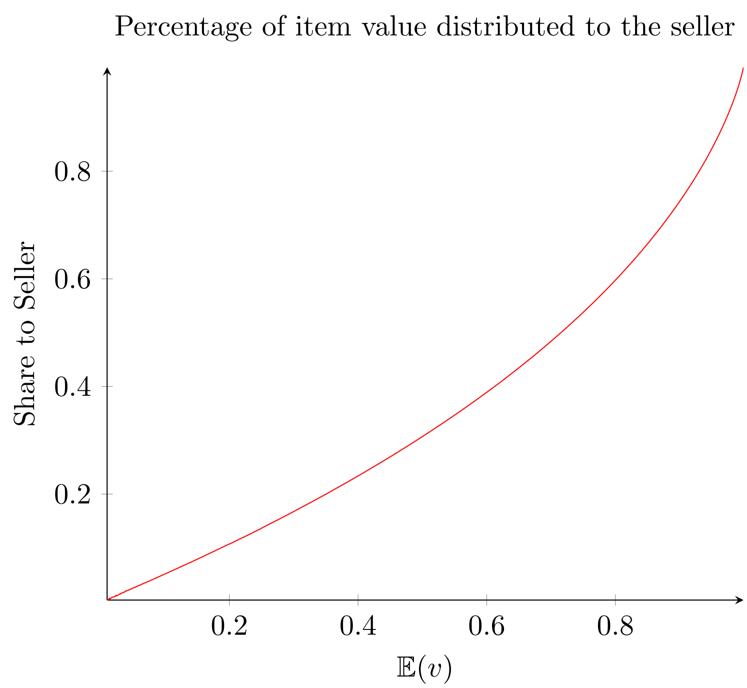

5.3 Binary Finite N Results

Now let us analyze a finite N binary distribution problem. Similar to the continuous distribution analysis, when we relax from asymptotic behavior, we lose the tractability of a simple solution and instead have a set of characterization equations. The equilibrium constraint is

$$\frac{\mathbb{E}(v) - R}{N} \geq (1 - \mathbb{E}(v))(0 - \delta) + \int \mathbb{1}(\delta > a)(1 - \delta) d\sigma(1, a).$$

Again, let $F(a)$ be the cdf of a given $v = 1$, so

$$F(\delta) = Prob(\delta > a | v = 1) = \frac{\int \mathbb{1}(\delta > a) d\sigma(1, a)}{\mathbb{E}(v)},$$

and the equilibrium constraint becomes

$$F(\delta) \leq \frac{\frac{\mathbb{E}(v) - R}{N} + \delta(1 - \mathbb{E}(v))}{(1 - \delta)\mathbb{E}(v)}.$$

To minimize revenue, we want the highest F as much as possible, so that every constraint is binding. When $\delta = 0$, to make sure the constraint is effective, it must be that

$$F(0) = \frac{\mathbb{E}(v) - R}{N}.$$

And as δ increases, F increases, so the cutoff $\overline{a}$ that satisfies $F(\delta) = 1$ for all $\delta \geq \overline{a}$ is given by the equation

$$\frac{\frac{\mathbb{E}(v)-R}{N} + \overline{a}(1 - \mathbb{E}(v))}{(1 - \overline{a})\mathbb{E}(v)} = 1. \tag{6}$$

Therefore, the distribution of a given $v = 1$ is

$$f(a) = F'(a) = \frac{1 - \mathbb{E}(v) + \frac{\mathbb{E}(v)-R}{N}}{(1 - a)^2\mathbb{E}(v)}, \text{ where } 0 < a \leq \overline{a},$$

and

$$\begin{aligned} R &= \int a d\sigma(1, a) \\ &= \int a f(a)\mathbb{E}(v) da \\ &= \int_0^{\overline{a}} (1 - \mathbb{E}(v) + \frac{\mathbb{E}(v) - R}{N}) \frac{a}{(1 - a)^2} da \\ &= \left[1 - \mathbb{E}(v) + \frac{\mathbb{E}(v) - R}{N}\right] \times \left(\frac{\overline{a}}{1 - \overline{a}} - \ln(1 - \overline{a})\right). \end{aligned} \tag{7}$$

Using equations (6) and (7), we could solve $\overline{a}$ and R, and R gives the binary distribution revenue guarantee with finite N buyers.

6 Beyond FPA and Common Values

In this section, we claim that all the analysis can be extended beyond the first-price auction to any standard auctions with no losing payments, and beyond common values to symmetric priors of the item values, and all characterization results will still apply.

In the proofs, we use the first-price auction structure to show that the first max must be equal to the second max. It is necessary to have the first-price auction design to derive this equality. However, for the following argument that it is always possible to create an identical-play minimum revenue BCCE, the first-price auction is no longer needed. We can replace it with any standard auction (where the highest buyers always win the item) in which all losers don't need to pay anything, and then all the results still follow. The idea of constructing a low revenue BCCE is to somehow distribute the positive payoff of the item winner to all other losing buyers: it reduces all other buyers' incentive to offer a higher price and get the item entirely. A possible way of doing so is to divide the positive surplus equally among all buyers (so they are homogeneous), and it achieves the worst-case revenue guarantee level as it fully exploits the equilibrium constraint. It is necessary to point out that under standard auctions with no losing payments, the revenue guarantee only holds over all BCCEs with a positive payoff for any buyer in the equilibrium. A positive payoff is needed to eliminate trivial zero-revenue scenarios. For example, in the second-price auction, it is possible that one buyer always aims for an arbitrarily large price and all others bid 0, and it is an uninteresting equilibrium with zero revenue.

In addition, the common value assumption can be relaxed to any symmetric prior belief of the buyers. Now the value $v_i \in \mu(v)$, and v_i can be idiosyncratic across buyer i. The only assumption here is the symmetric common prior $\mu(v)$. Under this generalized problem, we can still treat the value as if the common value. That is, for all buyer i, let $s_i = v$ for a common value v drawing from μ, and let $a_i = H(s_i)$ as defined before, then it satisfies all the constraints and attains the same revenue guarantee level as the common value case. The intuition is that a specific realized value v_i is unimportant, as the BCCE equilibrium deviation is ex-ante for a fixed action a'_i. Everything that matters is in the sense of the prior distribution of values, and symmetric priors make it possible to construct identical plays and let all buyers share surplus.

7 Conclusion

In this paper, under the notion of BCCE, we characterize the revenue guarantee when the value takes a general continuous distribution or a binary distribution based on an identical play simplification. We start with the first-price auction and common values but then show that it works for standard auctions and symmetric prior as well. Moreover, the asymptotic revenue guarantee of a continuous distribution is equal to that under BCE in the literature, and the binary distribution revenue guarantee gives the revenue guarantee level when the distribution is unknown to the seller except for the expected value. This paper enlarges our understanding of information-robust predictions beyond the scope of BCE and offers some tractable conclusions when BCCE is involved. Aside from theoretic values, it offers insights on the applied side: in many real-world Bayesian settings, when selling an item, the item owner has little knowledge, or even no clue, of the information structure held by the buyers. In such scenarios, it might be reasonable for the seller to target a safer but guaranteed level of return, which is heavily researched in this paper. Based on discussions in this paper, the supplier may now have some confidence in the level of expected revenues.

References

Azacis, Helmuts, and Peter Vida. 2015. "Collusive communication schemes in a first-price auction." *Economic Theory*, 58(1): 125–160.

Battigalli, Pierpaolo, and Marciano Siniscalchi. 2003. "Rationalizable bidding in first-price auctions." *Games and Economic Behavior*, 45(1): 38–72. First World Congress of the Game Theory Society.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2017. "First-Price Auctions With General Information Structures: Implications for Bidding and Revenue." *Econometrica*, 85(1): 107–143.

Blum, Avrim, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth. 2008. "Regret minimization and the price of total anarchy." 373–382.

Brooks, Benjamin, and Songzi Du. 2021. "Optimal Auction Design With Common Values: An Informationally Robust Approach." *Econometrica*, 89(3): 1313–1360.

Dekel, Eddie, and Asher Wolinsky. 2003. "Rationalizable outcomes of large private-value first-price discrete auctions." *Games and Economic Behavior*, 43(2): 175–188.

Du, Songzi. 2018. "Robust Mechanisms Under Common Valuation." *Econometrica*, 86(5): 1569–1588.

Fang, Hanming, and Stephen Morris. 2006. "Multidimensional private value auctions." *Journal of Economic Theory*, 126(1): 1–30.

Feldman, Michal, Brendan Lucier, and Noam Nisan. 2016. "Correlated and Coarse Equilibria of Single-Item Auctions." 131–144. Berlin, Heidelberg:Springer Berlin Heidelberg.

Hartline, Jason D, Vasilis Syrgkanis, and Éva Tardos. 2015. "No-regret learning in Bayesian games." *Advances in Neural Information Processing Systems*, 2015-January: 3061–3069.

Hart, Sergiu, and Andreu Mas-Colell. 2000. "A Simple Adaptive Procedure Leading to Correlated Equilibrium." *Econometrica*, 68(5): 1127–1150.

Jin, Yaonan, and Pinyan Lu. 2023. "First Price Auction is 1-1/e2 Efficient." *J. ACM*, 70(5).

Milgrom, Paul R., and Robert J. Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica*, 50(5): 1089–1122.

Moulin, Hervé, and Jean-Philippe Vial. 1978. "Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon." *International Journal of Game Theory*, 7: 201–221.

Syrgkanis, Vasilis, and Eva Tardos. 2013. "Composable and efficient mechanisms." *STOC '13*, 211–220. New York, NY, USA:Association for Computing Machinery.

Vickrey, William. 1961. "Counterspeculation, Auctions, And Competitive Sealed Tenders." *The Journal of Finance*, 16(1): 8–37.

A Proofs for Section 3 (Identical Plays)

Proposition 2. *Let the first max (the highest number in a set) be* $\max_1$*, the second max (the second highest number in a set) be* $\max_2$*, then under FPA, if* σ *is a min-R BCCE, for any* $(v, a_1, a_2, ..., a_n)$ *such that* $\sigma(v, a_1, a_2, ...a_n) > 0$*, it must be that* $\max_1\{a_1, ..., a_n\} = \max_2\{a_1, ..., a_n\}$

Proof. We apply a similar analysis approach as the case of 2 buyers but generalize it to N buyers. Define $\tilde{\sigma_j}$ be a permutation of σ by j indexes (i.e., $\forall i > 0$, buyer $i + j (\text{mod } n)$ in $\tilde{\sigma_j}$ behaves exactly the same as buyer i in σ.) For better notation purpose, we say a^{+j} is a permutation of $+j$ for the action profile a where $a = (a_1, a_2, ..., a_n)$, then our design can be rewritten as $\tilde{\sigma_j}(v, a^{+j}) = \sigma(v, a)$ for all v, j, and a. By the nature of permutations, $\forall j$, $\tilde{\sigma_j}$ must be a min-R BCCE.

Now by the convexity of BCCE and the linearity of revenue,

$$\widehat{\sigma} := \frac{1}{n}\sum_{j=1}^{n}\tilde{\sigma_j}$$

must also be a min-R BCCE. We will work on $\widehat{\sigma}$ from now on, as it gives advantages of identical plays and the same payoff levels for all the buyers. All the proof steps so far are analogous to the case of 2 buyers, and our next step is to modify $\widehat{\sigma}$, construct a new distribution γ, and, if the claim doesn't hold, draw contradictions from there.

For any $(v, a_1, a_2, ..., a_n)$ such that $\sigma(v, a_1, a_2, ...a_n) > 0$, let us assume that $\sigma(v, a_1, a_2, ...a_n) = p > 0$, so $\forall j,\ \tilde{\sigma_j}(v, a_{+j}) = \sigma(v, a) = p$. Define

$$\widehat{\sigma}(v) = \frac{1}{n}\sum_{j}\tilde{\sigma_j}(v, a^{+j}) \times (v, a^{+j}) = \frac{p}{n} \times \sum_{j}(v, a^{+j})$$

and let $a_{\max_1}$ be the first max of $\{a_1, a_2, ..., a_n\}$, and $a_{\max_2}$ be the second max of $\{a_1, a_2, ..., a_n\}$. Without loss of generality, we can let a_1 be the first max, and a_2 be the second max. Now consider a modified distribution

$$\begin{aligned}\gamma &= \widehat{\sigma} - \widehat{\sigma}(v) + \frac{(n-1)p}{n}(v, a_{\max_1}, a_{\max_1}, ..., a_{\max_1}) + \frac{p}{n} \times (v, a_{\max_2}, a_{\max_2}, ..., a_{\max_2)}\\ &= \widehat{\sigma} - \widehat{\sigma}(v) + \frac{(n-1)p}{n}(v, a_1, a_1, ..., a_1) + \frac{p}{n} \times (v, a_2, a_2, ..., a_2)\end{aligned}$$

In other words, we modify the dictionary $\widehat{\sigma}$ and construct a new dictionary γ in the following way: first delete the $\widehat{\sigma}(v)$ parts of $\widehat{\sigma}$ (subtract probability value p/n in the keys of (v, a^{+j}) for every j, then add probability value $(n-1)p/n$ to key $(v, a_1, a_1, ..., a_1)$ and probability value p/n $(v, a_2, a_2, ..., a_2)$. It is easy to check conditions that $\sum_a \sigma(v, a) = \mu(v)$ for all possible v and that $\sum_{v,a} \sigma(v, a) = 1$ are satisfied, so this modification is valid.

Now we want to compare γ to $\widehat{\sigma}$. First we show that $R(\widehat{\sigma}) \geq R(\gamma)$. By the nature of FPA pricing function,

$$R(\widehat{\sigma}) - R(\gamma) = \frac{p}{n} n a_{max_1} - \frac{p}{n}(n-1)a_1 - \frac{p}{n}a_2 = \frac{p}{n}(a_1 - a_2) \geq 0.$$

Consequently, since buyers' actions are both identical in $\widehat{\sigma}$ and γ, then $\forall i$,

$$E_{\widehat{\sigma}}(u_i) = \frac{EV - R(\widehat{\sigma})}{n} \leq \frac{EV - R(\widehat{\gamma})}{n} = E_\gamma(u_i)$$

Secondly, we show that γ is also a BCCE. Let $\widehat{\sigma}_{rest} = \widehat{\sigma} - \widehat{\sigma}(v)$, the common part shared in $\widehat{\sigma}$ and γ. Then from the perspective of BCCE equilibrium constraint, since $\widehat{\sigma}$ is a valid BCCE, (we use a different notation b as a is fixed in the previous argument),

$$\forall i,\ \forall b'_i, \sum_{v,b_i,b_{-i}} u_i(v, b_i, b_{-i})\widehat{\sigma}(v, b_i, b_{-i}) \geq \sum_{v,b_i,b_{-i}} u_i(v, b'_i, b_{-i})\widehat{\sigma}(v, b_i, b_{-i}).$$

Rewrite the equation,

$$\begin{aligned}\forall i,\ \forall b'_i, E_{\widehat{\sigma}}(u_i) &\geq \sum_{v,b_i,b_{-i}} u_i(v, b'_i, b_{-i})\widehat{\sigma}(v, b_i, b_{-i}) \\ &= \sum_{v,b_i,b_{-i}} u_i(v, b'_i, b_{-i})\widehat{\sigma}_{rest}(v, b_i, b_{-i}) + \frac{p}{n} \times \sum_j u_i(v, b'_i, a^{+j}_{-i})\end{aligned}$$

But if we consider the marginal opponent's action distribution, by the allocation rule of FPA, only the highest bid among all opponents $(-i)$ would affect buyer i's payoff and correspondingly, optimal strategies. Therefore, with a litter abuse of notation which restricts attention only to the highest bid of all opponents,

$$\sum_j u_i(v, b'_i, a_{-i}^{+j}) = \sum_j u_i(v, b'_i, \max(a_{-i}^{+j})) = (n-1)u_i(v, b'_i, a_1) + u_i(v, b'_i, a_2)$$
$$= (n-1)u_i(v, b'_i, a_1, a_1, ..., a_1) + u_i(v, b'_i, a_2, a_2, ..., a_2).$$

Then $\forall i,\ \forall b'_i$,

$$E_\gamma(u_i) \geq E_{\widehat{\sigma}}(u_i)$$
$$\geq \sum_{v, b_i, b_{-i}} u_i(v, b'_i, b_{-i})\widehat{\sigma}_{rest}(v, b_i, b_{-i}) + \frac{p}{n}[(n-1)u_i(v, b'_i, a_1, a_1, ..., a_1) + u_i(v, b'_i, a_2, a_2, ..., a_2)]$$
$$= \sum_{v, b_i, b_{-i}} u_i(v, b'_i, b_{-i})\gamma(v, b_i, b_{-i}).$$

Hence, γ is also a valid BCCE. However, because $\widehat{\sigma}$ is a minimum revenue BCCE, it must be that $R(\widehat{\sigma}) \leq R(\gamma)$. This is only allowed to occur when $\frac{p}{n}(a_1 - a_2) = 0$, or when $a_1 = a_2$. Hence, it must be that $\max_1 \{a_1, a_2, ..., a_n\} = \max_2 \{a_1, a_2, ..., a_n\}$, as desired. □

Lemma 1. *(Identical plays) There exists a min-R BCCE σ such that for all $\sigma(v, a_1, a_2, ..., a_n) > 0$, $\exists \alpha$ s.t. $a_1 = a_2 = ... = a_n = \alpha$. We call such σ an identical play BCCE.*

Proof. To prove this claim, we start with an arbitrary min-R BCCE σ and construct an identical play min-R BCCE γ based on it. First follow the procedure in proposition 2, as long as there is any tuple $(v, a_1, a_2, ..., a_n)$ such that $\sigma(v, a_1, a_2, ..., a_n) = p > 0$ and $a_1, a_2, ..., a_n$ are not all the same, fix $(v, a_1, a_2, ..., a_n)$, we can construct a min-R BCCE γ given σ. Notice that the first max must have the same value as the second max, so γ can be simplified as

$$\gamma = \widehat{\sigma} - \widehat{\sigma}(v) + \frac{np}{n}(v, a_{\max_1}, a_{\max_1}, ..., a_{\max_1})$$
$$= \widehat{\sigma} - \widehat{\sigma}(v) + p \times (v, a_1, a_1, ..., a_1)$$

By the subsequent argument in proposition 2, γ is a min-R BCCE. Hence, in this iteration of modification, we manage to create a min-R BCCE which achieves, for the fixed tuple $(v, a_1, a_2, ..., a_n)$, $a_1 = a_2 = ... = a_n = \alpha$ for $\alpha = a_1$. Using the newly created BCCE as the next start point, we could repeat iterations as long as there is any tuple $(v, a_1, a_2, ..., a_n)$ such that $\sigma(v, a_1, a_2, ..., a_n) = p > 0$ and $a_1, a_2, ..., a_n$ are not all the same. Eventually, we could find an identical play min-R BCCE, such that for all $\sigma(v, a_1, a_2, ..., a_n) > 0$, there exists some α such that $a_1 = a_2 = ... = a_n = \alpha$. □

B Proofs for Section 4 (Characterization of a Continuous Distribution)

Lemma 2. *(Monotonicity) Without loss of generality, we can treat the identical play min-R BCCE σ as monotonic, that is $\forall(v_1, a_1), (v_2, a_2)$ such that $\sigma(v_1, a_1) > 0, \sigma(v_2, a_2) > 0$, if $v_1 > v_2$, then $a_1 > a_2$.*

Proof. We start with an identical play min-R BCCE σ and pick a small share of the distribution with the probability ϵ in the regions (v_1, a_1) and (v_2, a_2) such that $\sigma(v_1, a_1) > 0$, $\sigma(v_2, a_2) > 0$, $v_1 > v_2$, and $a_1 < a_2$. Then we switch the actions in those 2 regions with ϵ probability and construct a new distribution $\tilde{\sigma}$, that is,

$$\tilde{\sigma} = \sigma - \epsilon \times (v_1, a_1) - \epsilon \times (v_2, a_2) + \epsilon \times (v_1, a_2) + \epsilon \times (v_2, a_1).$$

It is easy to check that after the changes, $\tilde{\sigma}$ is still an identical play min-R BCCE. The economic reason behind it is that previously, buyers paid less for a higher-valued item and paid more for a lower-valued item, so they would be more inclined to deviate, and now it is the other way around, so all the equilibrium constraints have to be satisfied. Then we could repeat the ϵ changes on all inconsistent pairs of (v_1, a_1) and (v_2, a_2). At the end of the day, a monotonic identical play min-R BCCE σ could be obtained. □

C Proofs for Section 5 (Unknown Distributions)

Proposition 3. *(Lower bound for binary restriction) Under FPA,*

$$\min_{\mu(v) \;\; s.t. E_\mu(v) = \mathbb{E}(v)} \min_{\gamma} R(\gamma)|\mu \Longleftrightarrow \min_{\gamma} R(\gamma)|\underline{\mu}$$

where γ *is a BCCE that satisfies* $\sum_a \gamma(v, a) = \mu(v)$ *for all* v*, and* $\underline{\mu}$ *is a binary distribution of* v *where* $\underline{\mu}(0) = 1 - \mathbb{E}(v)$*, and* $\underline{\mu}(1) = \mathbb{E}(v)$*. Note that here* $\mathbb{E}_{\underline{\mu}}(v) = \mathbb{E}(v)$*.*

Proof. Let σ be a minimum revenue BCCE across all μ with $\mathbb{E}_\mu(v) = \mathbb{E}(v)$, i.e.,

$$\sigma \in \arg\min_{\gamma} \left(\min_{\mu \ s.t. \mathbb{E}_\mu(v) = \mathbb{E}(v)} R(\gamma)|\mu \right).$$

Define μ_{min} to be such a value distribution when the minimum revenue is attained. Then by definition of minimum revenue BCCE across all distributions, for any $\tilde{\sigma}$ that is a BCCE following a value distribution $\tilde{\mu}$,

$$R(\sigma)|\mu_{min} \leq R(\tilde{\sigma})|\tilde{\mu}$$

Now consider a support in σ where $\sigma(v, a_1, a_2, ..., a_n) = p > 0$ (here v follows the μ_{min} distribution). By Proposition 4, we can choose σ to be an identical play BCCE such that $a_1 = a_2 = ... = a_n = c \leq v$. Then let

$$\tilde{\sigma} = \sigma - p \times \sigma(v, c, c, ..., c) + p(1 - v) \times (0, c, c, ..., c) + pv \times (1, c, c, ..., c)$$

In another sense, we modify the distribution σ and construct a new distribution $\tilde{\sigma}$ (under a different $\tilde{\mu}$) but still guarantee that they have the same expected value $\mathbb{E}(v)$ and share the same action profiles. Therefore, they have the same deviation payoff for any deviation d_i for a buyer i. Also because of same action profiles, $R(\tilde{\sigma})|\tilde{\mu} = R(\sigma)|\mu_{\min}$, and we can further conclude that $\tilde{\sigma}$ is also a min-R BCCE. Note that after the modification, we successfully replace the part of the distribution with value v to parts with the value of 0 and 1, and still keep the minimum revenue potential. We could repeat the procedure until there is no positive probability in values other than 0 and 1 (i.e., until $\tilde{\mu}(v) = 0$ for all $v \neq 0, 1$.) But since $\mathbb{E}_{\tilde{\mu}}(v) = \mathbb{E}(v)$, so it must be that $\tilde{\mu} = \underline{\mu}$. Hence, we can restrict our attention only to the binary $\underline{\mu}$ distribution. □

Proposition 4. *If* σ *is a min-R BCCE such that* $\forall (v, a_1, a_2, ..., a_n)$ *where* $\sigma(v, a_1, a_2, ..., a_n) > 0$*,* $a_1 = a_2 = ... = a_n = \alpha$ *for some* α*, then* $\alpha \leq v$*. Moreover, there exists a min-R BCCE such that* $\forall (v, a_1, a_2, ..., a_n)$ *where* $\sigma(v, a_1, a_2, ..., a_n) > 0$*,* $a_1 = a_2 = ... = a_n = \alpha \leq v$ *for some* α*.*

Proof. Start with σ, fix $(v, a_1, a_2, ..., a_n)$ and suppose that $\sigma(v, a_1, a_2, ..., a_n) = p > 0$, then $a_1 = a_2 = ... = a_n = \alpha$ for some number α, and it is sufficient to prove $\alpha \leq v$. If $\alpha > v$, let $\tilde{\sigma} = \sigma - p \times (v, \alpha, \alpha, ..., \alpha) + p \times (v, v, v, .., v)$, so $\tilde{\sigma}$ is the same distribution as σ except for p probability share of $(v, \alpha, \alpha, ..., \alpha)$ substituted by $(v, v, v, ..., v)$. By design, $\tilde{\sigma}$ is a valid distribution. Now consider a deviation d_i for some buyer $i > 0$. Since all buyers have identical strategies and thus identical payoffs, it is sufficient to only look at one buyer i.

If $d_i \leq v$ or $d_i \geq \alpha$, deviation d_i would get exactly the same payoff for buyer i in either σ or $\tilde{\sigma}$. Since σ is a BCCE, no d_i is favorable in σ. Also note that $\mathbb{E}_{\tilde{\sigma}}(u_i) > \mathbb{E}_{\sigma}(u_i)$ because in $\tilde{\sigma}$ buyers collude and pay less when the value is v. Combining these two findings we can conclude that d_i is also not preferred in $\tilde{\sigma}$.

If $v < d_i < \alpha$, we can consider the difference in payoffs for buyer i in $\tilde{\sigma}$ and σ. By deviating to d_i, buyer i will win the bid and get $(v - d_i)p < 0$ in $\tilde{\sigma}$. On the other hand, playing d_i will lose the bid and get a zero payoff in σ. Again, recall that $\mathbb{E}_{\tilde{\sigma}}(u_i) > \mathbb{E}_{\sigma}(u_i)$. Hence, if d_i is not preferred in σ, it will be even less profitable in $\tilde{\sigma}$.

Now we can conclude that $\tilde{\sigma}$ is also a BCCE. However, notice that $\sum_i \mathbb{E}_{\tilde{\sigma}}(u_i) > \sum_i \mathbb{E}_{\sigma}(u_i)$, so $R(\tilde{\sigma}) < R(\sigma)$, which contradicts with the min-R assumption of σ. Hence, it can't be that $\alpha > v$, which completes the proof of the first part of the claim. The second part is directly implied by the first part and Lemma 1.

□