

Coarse Revenue Guarantee in First Price Auction and Beyond*

Linchen Zhang[†]

February 15, 2026

Abstract

I analyze the **coarse revenue guarantee**, defined as the minimum expected revenue across all Bayes Coarse Correlated Equilibria (BCCE), in the First Price Auction (FPA). This approach generalizes the standard **revenue guarantee**—based on Bayes Correlated Equilibrium (BCE)—established by Bergemann, Brooks and Morris (2017) in the FPA. Because BCCE characterizes the limit points of no-regret learning dynamics, this metric provides a more robust lower bound for markets populated by algorithmic or AI agents. I characterize the coarse revenue guarantee using a reduction to “identical play equilibria,” in which bidders adopt identical strategies regardless of their private signals. I prove that while the coarse revenue guarantee is strictly lower than the BCE-based guarantee for any finite number of buyers, the two measures converge asymptotically as the market grows. Furthermore, I develop a coarse revenue guarantee ranking over standard auctions, generalizing the logic of Bergemann, Brooks and Morris (2019). I also extend the coarse analysis to all value distributions with a fixed expected value.

JEL Classification: D44, D82, C72

Keywords: Robust Mechanism Design, First Price Auction, Bayes Coarse Correlated Equilibrium, Revenue Guarantee

*I am indebted to Songzi Du for his invaluable guidance on this project. I am immensely grateful to Aram Grigoryan, Jun Liu, Joel Sobel, and Emanuel Vespa for their useful suggestions for this paper. I also thank Cuimin Ba, Arjada Bardhi, Aislinn Bohren, Benjamin Brooks, Daniel Gottlieb, Duarte Gonçalves, Menglong Guan, Jiangtao Li, Markus Moller, Joao Ramos, Doron Ravid, Denis Shishkin, Vasiliki Skreta, Alex Smolin, Joel Watson, Frank Yang, and Mengxi Zhang, as well as the audiences at EC 24', GAIMSS 24', Stony Brook Game Theory Conference 2024, and UCSD TBE Workshop for valuable conversations and constructive comments. A previous version of this paper was circulated as “Revenue Guarantee across Bayes Coarse Correlated Equilibria.” All errors are my own.

[†]Department of Economics, University of California San Diego. Email: liz032@ucsd.edu.

1 Introduction

The robustness of auction mechanisms to information structure is a central concern in modern mechanism design. In their seminal work, Bergemann, Brooks and Morris (2017) characterize the **revenue guarantee**—defined as the minimum expected revenue across all information structures consistent with a Bayes Correlated Equilibrium (BCE)—in the First Price Auction (FPA). Their analysis highlights that variations in buyers’ information can significantly impact revenue, even when the auction format and value distributions are fixed. However, the BCE framework relies on the assumption that agents optimize deviations conditional on their private signals (interim rationality). In many modern algorithmic and high-frequency trading environments, agents—whether human or artificial—employ heuristic learning protocols that may not satisfy strict conditional obedience. Instead, these dynamics typically converge to the broader set of Bayes Coarse Correlated Equilibria (BCCE) (Hartline, Syrgkanis and Tardos, 2015).

With the exponential growth of artificial intelligence in economic decision-making, markets are increasingly populated by automated agents trained on no-regret learning algorithms. Since the limit points of such dynamics are well-captured by BCCE rather than BCE, a worst-case revenue analysis based on BCCE provides a strictly more robust lower bound for the seller in these environments. In this paper, I define the **coarse revenue guarantee** as the minimum expected revenue across all BCCEs. I extend the robust mechanism design literature in two distinct dimensions: first, by generalizing the FPA analysis from the revenue guarantee (over BCE) to the coarse revenue guarantee (over BCCE); and second, by establishing a “Coarse Revenue Guarantee Ranking” over standard auctions, extending the insights of Bergemann, Brooks and Morris (2019).

My first contribution characterizes the seller’s worst-case revenue when buyers are rational only in the coarse sense: they have no incentive to deviate to a fixed action *ex-ante*, though they may experience *ex-post* regret. I demonstrate that for any finite number of buyers, the coarse revenue guarantee is strictly lower than the revenue guarantee characterized by Bergemann, Brooks and Morris (2017). This implies that in markets dominated by no-regret learners rather than sophisticated Bayesian updaters, the seller’s potential downside is more severe than previously understood. While it is intuitive that expanding the minimization domain from BCE to

BCCE lowers the guarantee, the technical contribution lies in the characterization. I introduce a reduction demonstrating that it suffices to examine “**identical play equilibria**,” in which all buyers take the same action regardless of their signals, to identify the worst case. This contrasts with Bergemann, Brooks and Morris (2017), where the worst-case BCE is achieved under an information structure where bidders know only whether they hold the highest value and the average value of $N - 1$ lowest competitors. While the BBM structure implies symmetric play conditional on signals, my analysis shows that under coarse rationality, the lower bound is driven by unconditional identical play. However, I prove that this “coarseness gap” vanishes in the limit: as the number of buyers N grows large, the information structure in BBM converges to the value distribution, and the asymptotic coarse revenue guarantee converges exactly to the standard revenue guarantee.

My second contribution extends the “Revenue Guarantee Ranking” and “Revenue Guarantee Equivalence” results of Bergemann, Brooks and Morris (2019) to the coarse setting. I find that the First Price Auction maximizes the coarse revenue guarantee among standard auctions. Furthermore, when restricting the analysis to identical play equilibria, all standard auctions yield an equivalent coarse revenue guarantee, which matches the bound derived for the FPA.

Beyond information uncertainty, this paper also addresses scenarios where the seller faces ambiguity regarding the shape of the value distribution. I prove that for a fixed expected value, the binary distribution (taking values 0 and 1) minimizes the coarse revenue guarantee. This result provides a uniform lower bound: if a seller knows only the mean of the asset’s value, the binary case represents the “worst-case prior,” providing a robust floor for revenue expectations.

The remainder of the paper is organized as follows. Section 2 discusses related literature. Section 3 introduces the model. Section 4 establishes the sufficiency of identical play equilibria. Section 5 characterizes the coarse revenue guarantee for continuous distributions and establishes asymptotic convergence to the Bergemann, Brooks and Morris (2017) benchmark. Section 6 states analogous results of revenue guarantee ranking in a “coarse” sense. Section 7 presents a binary distribution lower bound for unknown shapes of the value distribution. Section 8 concludes. Omitted proofs are in the Appendix.

2 Related Literature

This research contributes to three primary strands of literature: first-price auctions (FPA), robust revenue guarantees, and the intersection of Bayes Coarse Correlated Equilibrium (BCCE) with no-regret learning dynamics. First, this work builds upon the extensive literature on the FPA. Early studies—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)—relied on specific assumptions regarding information structures. More recently, Bergemann, Brooks and Morris (2017) allowed for information variations but restricted their analysis to Bayes Correlated Equilibrium (BCE). My paper advances this analysis by characterizing the “coarse” revenue guarantee under the weaker notion of BCCE. Furthermore, the results presented here extend beyond the FPA; they preserve revenue guarantee rankings and equivalence results similar to Bergemann, Brooks and Morris (2019), but establish them within the coarse setting.

Second, this paper contributes to the literature on revenue guarantees (e.g., Bergemann, Brooks and Morris (2017), Du (2018), and Brooks and Du (2021)) by being among the first to investigate these guarantees outside the scope of BCE. While there is existing work studying guarantees in terms of welfare analysis—specifically the Price of Anarchy, such as Syrgkanis and Tardos (2013) and Jin and Lu (2023)—this paper focuses strictly on revenue rather than welfare outcomes.

Finally, this work contributes to the literature on BCCE. Originally proposed by Moulin and Vial (1978), BCCE has historically received limited attention in economics, largely because the concept was viewed as too permissive regarding fixed deviation assumptions. However, the growing prevalence of AI agents trained on no-regret learning dynamics has created a need for robust BCCE predictions. As noted by Hartline, Syrgkanis and Tardos (2015) and supported by learning literature such as Hart and Mas-Colell (2000) and Blum et al. (2008), no-regret dynamics converge to the set of BCCE. Consequently, our analysis offers critical theoretical insights into the outcomes of industrially applied algorithmic designs (see also Lomys and Magnolfi (2025)).

3 Model

3.1 General Setup

Consider a single-unit first-price auction with a set of risk-neutral buyers $\mathcal{I} = \{1, \dots, N\}$. The seller owns a good. The buyers have a pure common value for the good denoted $v \in V = [v, \bar{v}] \subseteq \mathbb{R}_+$, distributed according to a probability measure $\mu \in \Delta(V)$. While my primary analysis focuses on common values, I extend to symmetric priors where each buyer i has a private value $v_i \sim \mu_i = \mu$ in Section 5.3. As we shall see later, the symmetric prior is “as if” the common value by permuting buyers’ indexes. For brevity, I define the model with only pure common values, but all the results applies naturally to symmetric priors of value distributions. The inclusion of symmetric priors also aligns with the findings of Bergemann, Brooks and Morris (2017), which discuss revenue guarantees under BCE.

The action profile is denoted by $a = (a_1, \dots, a_N) \in A = \prod_i A_i \subseteq V^N$. In the FPA, the highest bidder wins; if there are ties, the item is allocated with equal probability among the tying winners. The winner pays their bid. Buyer i ’s utility given value v and action profile a is $u_i(v, a) = q_i(a)(v - a_i)$, where $q_i(a)$ is the allocation rule described above: it equals 1 if winning and 0 otherwise.

3.2 Equilibria and Revenue Guarantees

The buyers’ correlated information about the common value is described by an *information structure* $I = (S, \theta)$, where S_i is a finite set of signals for buyer i , $S = \prod_i S_i$, and $\theta \in \Delta(V \times S)$ is the joint distribution of the values and signals. We let $\mathcal{I}(\mu)$ be the set of information structures for which $\text{marg}_V \theta = \mu$.

Given the information structure I , a *strategy* for buyer i is a mapping $b_i : S_i \rightarrow \Delta(A)$, that assigns to each signal s_i a likelihood of bidding $b_i(\cdot | s_i)$. Buyer i ’s expected utility given a strategy profile b is

$$U_i(I, b) \equiv \int_{v, s, a} u_i(v, a) \left(\prod_i b_i(a_i | s_i) \right) \sigma(dv, ds).$$

A *Bayes correlated equilibrium* for given I is a recommended strategy profile b such that $U_i(I, b) \geq U_i(I, (b'_i, b_{-i}))$ for all strategy b'_i and all buyer i . Let $\mathcal{BCE}(I)$ be the set of BCE based on I .

On the other hand, a *Bayes coarse correlated equilibrium* given I is a recommended strategy profile b such that $U_i(I, b) \geq U_i(I, (b'_i, b_{-i}))$ for all strategy $b'_i = a'_i$ that does not depend on s_i and for all buyer i . Let $\mathcal{BCCE}(I)$ be the set of BCCE based on I .

Let $R(I, b)$ be the expected revenue at an equilibrium b :

$$R(I, b) \equiv \int_{v, s, a} \max_i a_i \left(\prod_i b_i(a_i | s_i) \right) \sigma(dv, ds).$$

Define the *revenue guarantee* under the prior μ as the infimum expected revenue over all information structures $I \in \mathcal{I}(\mu)$ and all equilibria $b \in \mathcal{BCCE}(I)$:

$$RG(\mu) \equiv \inf_{I \in \mathcal{I}(\mu)} \inf_{b \in \mathcal{BCCE}(I)} R(I, b).$$

Similarly, define the *coarse revenue guarantee* under the prior μ as the infimum expected revenue over all information structures $I \in \mathcal{I}(\mu)$ and all equilibria $b \in \mathcal{BCCCE}(I)$:

$$CRG(\mu) \equiv \inf_{I \in \mathcal{I}(\mu)} \inf_{b \in \mathcal{BCCCE}(p, I)} R(I, b).$$

3.3 BCCE and CRG, in shortened notation

Note that the solution concept of Bayes Coarse Correlated Equilibrium (BCCE) is robust to the underlying signals, as the deviation can only be a fixed action. In other words, signals are not particularly useful in the BCCEs. It is then convenient to define the equilibrium directly on the joint distribution of values and actions, $\sigma \in \Delta(V \times A)$, where $\text{marg}_V \sigma = \mu$.

Definition 1 (Bayes Coarse Correlated Equilibrium). *A joint distribution $\sigma \in \Delta(V \times A)$ is a BCCE if for all buyers i and all fixed deviations $a'_i \in A_i$:*

$$\int_{V \times A} u_i(v, a_i, a_{-i}) d\sigma(v, a) \geq \int_{V \times A} u_i(v, a'_i, a_{-i}) d\sigma(v, a).$$

This condition implies that no buyer can strictly gain by committing ex-ante to a specific action a'_i , effectively ignoring any private information. This contrasts with Bayes Correlated Equilibrium (BCE), which requires obedience conditional on every realized signal. As such, BCCE is a coarser and more inclusive concept than BCE.

Similarly, the coarse correlated equilibrium can be redefined as

Definition 2 (Coarse Revenue Guarantee).

$$CRG = \inf_{\sigma \in BCCE(\mu)} R(\sigma), \text{ where } R(\sigma) = \int_{V \times A} \max_i a_i d\sigma(v, a).$$

A BCCE that attains this infimum is defined as a min-R BCCE.

4 Identical Play Reduction

To tractably solve for the coarse revenue guarantee, I first reduce the dimensionality of the problem. I demonstrate that it is without loss of generality to restrict attention to *identical play equilibria*, where all buyers submit the same bid with probability one, regardless of signals.

Proposition 1 (Identical Plays). *There exists a min-R BCCE σ such that for all (v, a) in the support of σ , the actions are identical: $a_1 = a_2 = \dots = a_N$.*

Proof Sketch. I begin with the $N = 2$ case. If an asymmetric min-R BCCE σ exists, I can construct a symmetric equilibrium $\hat{\sigma}$ by averaging σ with its permutation. The convexity of the BCCE set ensures $\hat{\sigma}$ remains a min-R BCCE. I then construct a new distribution γ by “pooling” probability mass: if $\hat{\sigma}$ places weight on asymmetric bids (a_1, a_2) , γ redistributes this weight to (a_1, a_1) and (a_2, a_2) . All other parts of γ and $\hat{\sigma}$ are the same. I subsequently show that γ is still a BCCE but now with lower revenue than $\hat{\sigma}$. Consequently, it violates the σ min-R BCCE assumption. Hence, min-R BCCE must be identical plays with 2 buyers.

Alternatively, we can think through Jensen’s Inequality. In a First-Price Auction, the seller’s revenue is determined by the maximum bid. By Jensen’s inequality arguments applied to the max operator, revenue is minimized when bids are identical ($a_1 = a_2$), as price competition is eliminated while preserving incentive compatibility constraints.

For general N , I iteratively apply this symmetrization logic. I first establish that in any min-R BCCE, the highest and second-highest bids must be equal. This ensures the price is set by at least two buyers. Once the winning price is established, lower bids can be raised to match the winner without altering the allocation or payment, eventually converging to a fully identical profile $a_1 = \dots = a_N$. (See Appendix A for the detailed proofs regarding 2 buyers, N buyers, and the full reduction.) \square

This reduction allows me to treat the N buyers as a single representative agent in the worst case. The problem reduces to finding the joint distribution of the value v and a single joint bid a .

5 Characterization

I now characterize the coarse revenue guarantee for a continuous value distribution μ with PDF f and CDF F . I assume continuity to avoid heavier notation, but the results can be extended to discontinuous distributions.

5.1 Main Result

Theorem 1. *The Coarse Revenue Guarantee CRG and a threshold \hat{v} are uniquely determined by the system:*

$$\frac{\mathbb{E}[v] - CRG}{N} = \int_0^{\hat{v}} v f(v) dv, \quad (1)$$

$$CRG = \int_0^1 H(v) f(v) dv, \quad (2)$$

where the bidding strategy $H(v)$ for $v \geq \hat{v}$ is given by:

$$H(v) := \mathbb{E}[z \mid \hat{v} \leq z \leq v] \times \frac{F(v) - F(\hat{v})}{F(v)}. \quad (3)$$

Proof Intuition. The identical play reduction implies the equilibrium constraint must hold for any deviation δ : the representative buyer's equilibrium payoff must exceed the payoff from deviating to δ (winning all items where $a < \delta$). To minimize revenue, these constraints must bind for all relevant δ . I construct a monotonic equilibrium where bids a are strictly increasing in v above a threshold \hat{v} . Below \hat{v} , bids are zero. The threshold \hat{v} effectively balances the participation constraint (Equation 1). The function $H(v)$ is derived from the binding incentive constraints for $\delta > 0$. (See Appendix B for the detailed characterization and proofs.) \square

5.2 Asymptotic Convergence

An important implication of Theorem 1 is the relationship between the coarse revenue guarantee (BCCE) and the revenue guarantee (BCE) as the market grows.

Corollary 1. *As $N \rightarrow \infty$, the coarse revenue guarantee converges to the revenue guarantee in the first-price auction:*

$$\lim_{N \rightarrow \infty} CRG = \lim_{N \rightarrow \infty} RG = \int_0^1 G(v) f(v) dv,$$

where $G(v) = \mathbb{E}[z \mid z \leq v]$.

While the coarse revenue guarantee is strictly lower than the revenue guarantee for finite N (due to weaker obedience constraints in BCCE), this gap vanishes asymptotically. For example, with a Uniform $[0, 1]$ distribution and $N = 2$, $CRG \approx 0.047$, whereas $RG = 1/6 \approx 0.167$. However, both converge to $1/4$ as $N \rightarrow \infty$.

Alternatively, this asymptotic equivalence between the two lower bounds follows from the convergence of the two information structures in the worst case. The worst-case BCE in Bergemann, Brooks and Morris (2017) is achieved under an information structure in which bidders know only whether they hold the highest value and the average value of the $N - 1$ lowest bidders. As the number of buyers N grows large, the information structure in BBM converges to the value distribution, and the asymptotic coarse revenue guarantee converges exactly to the standard revenue guarantee. (See Appendix C for the detailed asymptotic characterization.)

5.3 Symmetric Priors

The assumption of common values can be relaxed to symmetric priors where the private value v_i are drawn from a common μ for every buyer i . I can construct the worst-case information structure by perfectly correlating buyers' signals so that $s_i = v$, thereby replicating the common-value environment. Essentially, the realization of private item values given the signals is irrelevant to agents using BCCE heuristics, since they only care about ex-ante optimization before receiving any signal. Therefore, as long as the distribution μ is the same for all buyers, we can treat it as if the "common value" case. However, as we show in Section 7, the shape of the distribution matters. This is also evident from equation 1: it affects the bidding strategy $H(v)$ in the minimum-revenue BCCE.

6 Coarse Revenue Guarantee Ranking

In this section, I establish a coarse revenue guarantee ranking and an equivalence result analogous to Bergemann, Brooks and Morris (2019), but adapted to the coarser solution concept of BCCE. Following their framework, I define standard auctions as those that (i) allocate the item to the highest bidder, and (ii) admit a BCCE in monotonic pure strategies under the symmetric independent private value model. This class naturally includes the first-price, second-price, and English auctions.

Consistent with the findings of Bergemann, Brooks and Morris (2019), the FPA achieves the highest coarse revenue guarantee among all standard auctions. This result stems from the universality of the worst-case configuration: the “identical play equilibrium” that characterizes the coarse revenue guarantee for the FPA constitutes a valid BCCE in any other standard auction. Since the worst-case outcome of the FPA is feasible in all standard auctions, the minimum expected revenue (guarantee) of any other standard auction cannot exceed that of the FPA.

Moreover, if we restrict our attention specifically to identical play equilibria, all standard auctions yield an equivalent lower bound on revenue. In an identical play equilibrium—where all bidders adopt the same strategy regardless of their private signals—the distinct pricing rules of various standard auctions effectively converge. Because the allocation relies purely on tie-breaking among identical bids, the strategic differences typically driven by bid shading or truthful revelation vanish. This result serves as a natural extension of the equivalence theorems in Bergemann, Brooks and Morris (2019) to the robust setting of BCCE.

7 Unknown Value Distribution

I next address the scenario in which the seller knows only the expected value $\mathbb{E}[v]$, but the distribution’s shape is unknown.

Proposition 2 (Binary Lower Bound). *For any distribution μ with expected value E , the coarse revenue guarantee is bounded below by the guarantee of a binary distribution $\underline{\mu}$ supported on $\{0, 1\}$ that has an expected value E :*

$$\min_{\mu: \mathbb{E}_{\mu}[v]=E} CRG(\mu) = CRG(\underline{\mu}).$$

Proof Sketch. Let σ be a min-R BCCE for a general distribution μ . I can apply a mean-preserving spread to the values while adjusting the probability weights to maintain the marginal distribution of actions. Specifically, probability mass at an interior v can be split to endpoints $\{0, 1\}$. This procedure preserves the expected value and the deviation incentives (which depend on the action distribution) but typically lowers the correlation between values and bids, reducing the seller’s revenue. Iterating this logic, the revenue is minimized when all probability mass is concentrated at the bounds 0 and 1.

(See Appendix D for the detailed proof.) □

For the binary distribution $\underline{\mu}$ with mean E , the asymptotic ($N \rightarrow \infty$) coarse revenue guarantee admits a closed-form solution:

$$CRG(\underline{\mu}) = E + (1 - E) \ln(1 - E). \tag{4}$$

This result implies that skewness favors the seller; a binary distribution (maximum variance) represents the worst-case scenario for revenue.

8 Conclusion

This paper characterizes the coarse revenue guarantee for auctions under the solution concept of Bayes Coarse Correlated Equilibrium. By identifying that worst-case outcomes are driven by identical play strategies, I provide a tractable characterization of coarse revenue guarantee in first-price auctions. I show that while the coarse revenue guarantee (BCCE) is strictly lower than the revenue guarantee (BCE) for finite markets, the two converge asymptotically. I extend the revenue guarantee ranking (Bergemann, Brooks and Morris (2019)) to a ranking of coarse revenue guarantee ranking among standard auctions. Furthermore, I establish the binary distribution as the universal lower bound for revenue given a fixed expected value. These results provide a comprehensive “worst-case” benchmark for sellers facing uncertainty about both the information structure and the strategic sophistication of buyers.

References

- Azacis, Helmut, 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.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris.** 2019. “Revenue Guarantee Equivalence.” *American Economic Review*, 109(5): 1911–29.
- Blum, Avrim, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth.** 2008. “Regret minimization and the price of total anarchy.” *Proceedings of the Annual ACM Symposium on Theory of Computing*, 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 $1-1/e^2$ Efficient.” *J. ACM*, 70(5).
- Lomys, Niccolò, and Lorenzo Magnolfi.** 2025. “Estimation of Games under No Regret: Structural Econometrics for AI.” *Working Paper*.
- 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 Appendix: Lemmas for Identical Play Reduction

This section establishes Proposition 1 (Identical Plays), which states that there exists a min-R BCCE of identical bids for all buyers. The proof proceeds in three steps, formalized by the following three lemmas. Lemma 1 handles the base case ($N = 2$), showing that symmetry is required for revenue minimization. Lemma 2 extends this logic to N buyers to establish that the top two bids must be equal. Finally, Lemma 3 uses this equality to construct a fully identical strategy profile, thereby completing the reduction.

Lemma 1 (Symmetry in 2-Buyer Case). *With $N = 2$ buyers, if σ is a min-R BCCE, then for any (v, a_1, a_2) in the support of σ , $a_1 = a_2$.*

Proof. Suppose σ is an asymmetric min-R BCCE. Construct the symmetrized distribution $\hat{\sigma} = \frac{1}{2}\sigma + \frac{1}{2}\tilde{\sigma}$, where $\tilde{\sigma}$ permutes buyer indices. By convexity of the BCCE set and linearity of revenue, $\hat{\sigma}$ is also a min-R BCCE. Suppose $\hat{\sigma}(v, a_1, a_2) = p > 0$ with $a_1 \neq a_2$. By symmetry, $\hat{\sigma}(v, a_2, a_1) = p$. Construct γ by redistributing this mass to symmetric profiles:

$$\gamma = \hat{\sigma} - p[(v, a_1, a_2) + (v, a_2, a_1)] + p[(v, a_1, a_1) + (v, a_2, a_2)].$$

We verify γ is a BCCE with lower revenue.

1. **Revenue:** In FPA, revenue is $\max(a_i)$. Since $a_1 \neq a_2$, $\max(a_1, a_2) > \frac{a_1 + a_2}{2}$. The revenue difference is $2p[\frac{a_1 + a_2}{2} - \max(a_1, a_2)] < 0$. Thus $R(\gamma) < R(\hat{\sigma})$.
2. **Incentives:** The deviation payoff depends on the marginal distribution of the opponent's bid. In $\hat{\sigma}$, conditional on this subset, the opponent bids a_1 or a_2 with equal probability. In γ , the opponent also bids a_1 or a_2 with equal probability. Thus, ex-ante incentives are unchanged.

This contradicts the minimality of σ . Thus $a_1 = a_2$. □

Lemma 2 (Equality of Highest Bids). *For any N , if σ is a min-R BCCE, then for any profile a in the support, the highest bid equals the second highest bid: $\max_1\{a\} = \max_2\{a\}$.*

Proof. Let $\hat{\sigma}$ be the symmetrization of σ . Suppose there exists a profile with unique winner $a_{(1)} > a_{(2)}$. We construct γ by shifting probability mass from the asymmetric

permutation orbit to the profile where all buyers bid $a_{(1)}$ (probability $\frac{N-1}{N}$) and all bid $a_{(2)}$ (probability $\frac{1}{N}$). The revenue decreases because the original revenue is $a_{(1)}$, while the new revenue is $\frac{N-1}{N}a_{(1)} + \frac{1}{N}a_{(2)} < a_{(1)}$. For incentives, consider the marginal distribution of the highest competing bid faced by buyer i . In the original equilibrium, a loser faces $a_{(1)}$, and the winner faces $a_{(2)}$. In γ , any buyer faces $a_{(1)}$ with probability $\frac{N-1}{N}$ and $a_{(2)}$ with probability $\frac{1}{N}$. Since the deviation utility is convex in the distribution of opposing bids (specifically, losing against higher bids reduces deviation payoff), the stronger competition in γ relaxes or maintains the obedience constraints. Thus γ is a valid BCCE with strictly lower revenue, a contradiction. \square

Lemma 3 (Construction of Identical Play). *There exists a min-R BCCE with identical play $a_1 = \dots = a_N$.*

Proof. Using the result from Lemma 2, the top two bids must be equal. This fixes the price. We can iteratively raise any losing bid $a_k < a_{(1)}$ to $a_{(1)}$ without affecting the allocation or price. This strictly tightens the obedience constraints for other players (who now face stiffer competition), ensuring the equilibrium conditions hold. Repeating this for all lower bids yields a profile where $a_1 = \dots = a_N$. \square

B Appendix: Proof for Characterization (Theorem 1)

This appendix provides the proof for Theorem 1. We first establish in Lemma 4 that the worst-case strategy must be monotonic, which allows us to define the inverse bidding function needed to solve the differential equation system characterizing the coarse revenue guarantee.

Lemma 4 (Monotonicity). *Without loss of generality, the identical play min-R BCCE σ is monotonic: $v_1 > v_2 \implies a(v_1) \geq a(v_2)$.*

Proof. Follows from standard rearrangement inequalities. Swapping actions to match the value order (a_{high} with v_{high}) maintains revenue but relaxes obedience constraints by reducing the incentive to deviate (buyers would rather deviate against low bids when they have high values; monotonicity prevents this). \square

Proof of Theorem 1

Treating buyers as a single agent with action $a(v)$ (justified by Proposition 1), the BCCE condition requires that for any deviation δ :

$$\frac{\mathbb{E}[v] - R}{N} \geq \int \mathbb{1}(\delta > a)(v - \delta) d\sigma(v, a). \quad (5)$$

Let $\bar{v}(\delta) = a^{-1}(\delta)$, which exists by Lemma 4. In a revenue-minimizing equilibrium, constraints must bind for all δ in the support. For small δ (and a), the participation constraint binds. Let \hat{v} be the threshold where bidding begins:

$$\frac{\mathbb{E}[v] - R}{N} = \int_0^{\hat{v}} v f(v) dv. \quad (6)$$

For $v > \hat{v}$, the constraint binds for all deviations. Differentiating the binding constraint integral equation yields the bidding function. Specifically, equating the payoff at deviation δ to the fixed equilibrium payoff:

$$\int_0^{\bar{v}(\delta)} (v - \delta) f(v) dv = \int_0^{\hat{v}} v f(v) dv.$$

Rearranging terms implies:

$$\delta F(\bar{v}(\delta)) = \int_{\hat{v}}^{\bar{v}(\delta)} v f(v) dv.$$

Solving for δ (which is $H(\bar{v})$) yields:

$$H(v) = \frac{1}{F(v)} \int_{\hat{v}}^v z f(z) dz = \mathbb{E}[z \mid \hat{v} \leq z \leq v] \frac{F(v) - F(\hat{v})}{F(v)}.$$

The revenue is $R = \int_{\hat{v}}^1 H(v) f(v) dv$. The system is closed.

C Appendix: Asymptotic Convergence (Corollary 1)

Proof of Corollary 1

As $N \rightarrow \infty$, the LHS of the obedience constraint $\frac{\mathbb{E}[v]-R}{N}$ vanishes. The constraint becomes:

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

For the constraint to bind everywhere (minimizing revenue), we require $\hat{v} \rightarrow 0$. The integral equation simplifies to:

$$\delta F(\bar{v}(\delta)) = \int_0^{\bar{v}(\delta)} v f(v) dv.$$

This implies $a(v) = \mathbb{E}[z \mid z \leq v]$, which matches the conditional expectation function $G(v)$ from Bergemann, Brooks and Morris (2017). Thus, $CRG \rightarrow \int_0^1 G(v) f(v) dv = RG$.

D Appendix: Binary Lower Bound (Proposition 2)

This section proves Proposition 2 (Binary Lower Bound). We first use Lemma 5 to restrict the domain of the optimization problem by showing bids cannot exceed values. We then proceed to show that the binary distribution minimizes revenue under these conditions.

Lemma 5 (Bid Bound). *If σ is a min- R BCCE with identical play α , then $\alpha \leq v$ almost surely.*

Proof. If $\alpha > v$, the buyer obtains negative utility. By deviating to 0, the buyer ensures 0 utility. To satisfy obedience, this loss must be compensated elsewhere. However, strictly lowering α to v improves buyer utility and lowers seller revenue, contradicting the minimality of R . \square

Proof of Proposition 2

Consider the binary distribution on $\{0, 1\}$ with mean E . At $v = 0$, $a = 0$. At $v = 1$, let the bid be distributed according to CDF F_a . The equilibrium constraint for deviation δ is:

$$0 \geq (1 - E)(0 - \delta) + E \int \mathbb{1}(\delta > a)(1 - \delta) dF_a(a).$$

Binding this constraint implies:

$$\Pr(a < \delta) = \frac{1 - E}{E} \frac{\delta}{1 - \delta}.$$

This defines the CDF that puts the maximum probability weight on low bids. The resulting revenue is:

$$R = E \int a dF_a(a) = E + (1 - E) \ln(1 - E).$$

For any general distribution μ , applying a mean-preserving spread to the boundaries $\{0, 1\}$ breaks the correlation between values and bids for interior points, forcing the seller to rely on the worst-case (boundary) obedience constraints. Thus the binary distribution yields the global lower bound.