Evidence from the Synthetic Laboratory: Language Models as Auction Participants

Anand Shah* MIT Sloan
Kehang Zhu* Harvard University
Yanchen Jiang Harvard University
Kerem Dayi Harvard University
Jeffrey Wang Harvard University
John J. Horton MIT Sloan
David C. Parkes Harvard University

Abstract

This paper investigates auction behavior of simulated AI agents (large language models, or LLMs). We begin by benchmarking these LLM-driven agents against established lab experiments across various auction settings: independent private value, affiliated private value, and common value auctions. Our findings reveal that LLM agents exhibit many behavioral traits similar to those observed in human participants within lab environments. Building on this, we investigate multi-unit combinatorial auctions under three distinct bid formats: simultaneous, sequential, and menu-based. Our study contributes fresh empirical insights into this classical auction framework. We run 1,000+ auctions for less than $100 with GPT-4, and develop a framework flexible enough to run auction experiments with any LLM model and a wide range of mechanism specification.

1 Introduction

The field of mechanism design, auction theory in particular, has benefited enormously from a rich interplay between empirics and theory. One recent example might be the development of *obviously strategy-proof mechanisms* (hereafter, OSP) (Li, 2017). The technical refinement was inspired by an empirical puzzle: despite it being well-known that the open-ascending clock (English) and second-price sealed-bid auctions were strategically equivalent, experiments since the 80s suggested that people were much 'better' at playing the open-ascending clock auction versus the sealed-bid auction Kagel et al. (1987). Motivated by empirical evidence, OSP provided one articulation for why the clock format might be 'better' than a sealed-bid format, and has since inspired a flourishing of work in auction design under behavioral constraints. The story echoes a well-understood but worth emphasizing point: empirical work is vital to the development of new theory.

Unfortunately, empirical evidence is quite expensive to generate. Li (2017)'s OSP experiments alone, with 404 participants, cost over $15,000.[1] The rise of LLMs raises the exciting new question as to whether there exist cheaper data generating processes that can substitute for human data for the purpose of studying human behavior, whether in economic systems or otherwise (Bubeck et al. (2023); Horton (2023); Manning et al. (2024)). The present work examines this question for auctions.

In particular, we provide evidence in three environments. We begin by considering classic 'revenue equivalence' results, comparing play under the first-price and second-price sealed-bid auctions in an independent, private value setting. We then proceed to make the auction 'easier' and 'harder' to play. First, we consider 'easier' auctions in the obviously strategy-proof sense: we compare play

[1] Li mentioned that for each participant, he would pay them $20 for participation and an additional money prize they won during the game. In total, he paid on average $37.47 for each participant.

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

between the second-price sealed-bid auction and two versions of the second-price ascending auction. We conduct these experiments in both affiliated private values and independent private values for robustness. Finally, we consider 'harder' auctions – we run the sealed-bid auctions in a common values setting to observe play when problems of adverse selection (i.e., the winner's curse) can bite. In all experiments, we observe monotone bidding and bid variation across values, suggesting that different plans can induce a distribution of LLM play. Our results are further compatible one common theme observed in the existing empirical literature: humans are risk-averse by nature, and this risk-aversion introduces a wedge in between expected theory and practice. LLM play in our experiments here is consistent with play we would expect from a risk averse set of humans. However, we emphasize that - because it's impossible to read the motivations of an LLM - we cannot disambiguate results that come from risk aversion and results that are driven by naive text completion that would not exist in out of distribution experiments.

To obtain the data for these empirical results, we have developed a code repository to systematically run experiments with some number of bidders and any prompting language. In particular, our repository is flexible enough that it can be used to generate synthetic data for almost any describable format with single or multiple goods.[2] For the experiments herein, we ran more than $1,000$ auctions with more than $5,000$ GPT-4 agent participants for costs totaling less than $\$100$. In contrast, the largest survey of auction experiments to date comes from Cox et al. (1988) of $1,500+$ auctions, with total costs likely considerably higher.

We acknowledge potential limitations in interpreting LLM responses as direct proxies for human responses, which should be approached with caution. Further experimental validation is needed, along with careful handling of data and specific use cases, to ensure reliability in practical applications.

2 Related work

LLMs as simulated agents: Recent LLMs, having been trained on an enormous corpus of human-generated data, are able to generate human-like texts and reason Achiam et al. (2023); Bubeck et al. (2023). Yet, they are far from perfect and show limited planning abilities and various cognitive biases Wan et al. (2023). There is a growing literature on using these human-like AI models as simulated agents in economics and social science studies Aher et al. (2023); Park et al. (2023); Brand et al. (2023). In this literature, Horton (2023) replicates four classical behavioral economics experiments by endowing a single LLM agent with different personas and querying it about its decisions and preferences.

LLMs in auctions: There are a few works on systematically using LLM as simulated agents in auction experiments. Fish et al. (2024) study the collusion behaviors in first-price sealed-bid auction of two LLM agents under the context of LLMs as a price setter for companies. Chen et al. (2023) study how to make an LLM better at playing auctions than humans. And Manning et al. (2024) ran a more limited study an a variant of an open-ascending clock auction with three bidders, focusing on deviations from rational economic theory in considering bidders' values and the final clearing price.

3 Benchmarks with previous auction experiments

3.1 FPSB vs. SPSB with IPV

We first consider the FPSB and SPSB auctions in IPV settings.

Setting: There are 3 bidders in each auction, and bidders draw an independent, private value from a uniform distribution $v \sim U[0, 99]$. Bidders, upon observing their value, submit a sealed bid. In the FPSB auction, the highest bidder pays her bid and receives the prize (and all other bidders pay 0 and receive no prize). In the SPSB auction, the highest bidder pays the second-highest bid and receives the prize (and all other bidders pay 0 and receive no prize). Bids are submitted in $\$1$ increments and ties are resolved randomly.

[2] We will make the code-base public soon and hope this will facilitate additional empirical work.

3.1.1 Theoretical and Empirical benchmarks

It is well-known that the SPSB auction has bidding one's value as a dominant strategy equilibrium. The FPSB auction, of course, has no equilibrium in dominant strategies but has a NE of bidding $\beta(v) = \frac{n-1}{n}v = \frac{2}{3}v$ as values are uniformly distributed with common support.

Experimental evidence for the FPSB persistently has bids above the risk-neutral NE prediction, suggesting a failure of revenue equivalence due to risk-aversion. Experimental data for the SPSB has agents bidding higher than in the FPSB, and sometimes even higher than their value. In both the FPSB and SPSB auction (and indeed, almost all auctions) there is robust evidence of bids being strictly monotone in one's value.

3.1.2 Simulation evidence

Simulations are run according to the setting above. Results are summarized in Figure 1.

Figure 1 demonstrates evidence of monotone bidding and the SPSB bids being larger than FPSB bids for the same value. However, there's fairly weak separation between the two bidding curves. There's also no bidding above one's value, which is a marked difference from the existing experimental evidence – usually, people find the inefficiency of bidding above one's value to be a subtle point in the SPSB auction. This may be an improvement of LLM play over human play.

Experiment logs in which LLMs explain their bidding decisions suggest that one reason to explain the weak separation in our data between the FPSB and SPSB auctions is that 1) LLMs are quite risk-averse and 2) that they sometimes confuse the SPSB and FPSB auctions. In this way, despite playing more intelligently than humans (in that LLMs almost never bid above their value), they do so because they may be confusing the SPSB for the FPSB. Additionally, LLMs in these classic settings demonstrate one new quirk not yet documented in experimental literature – bidding zero upon becoming frustrated. In 3% of runs, LLMs bid 0, often justifying their decision by arguing little chance of winning the good with a higher bid.

3.2 Obvious strategy-proofness

Next, we consider clock formats against sealed-bid formats. Following Li (2017) and Breitmoser & Schweighofer-Kodritsch (2022)'s experiments, we consider clock auctions against the SPSB auction in the affiliated private values (APV) setting.

Setting: Once again, there are 3 bidders in each auction but now bidders draw affiliated private values of the form $v = c + p$. The common component is drawn uniformly $c \sim U[0, 79]$ and the private component is drawn uniformly $p \sim U[0, 20]$. Winners of the auction receive their own value of the prize v when they win, so the 'common' and 'private' components only serve to make values correlated (even if draws are independent). The ascending clock auction (called AC below) is the classic English auction. The blind ascending clock auction (called AC-B below) is the English auction with the addition of not being told when other bidders leave. The SPSB auction was defined above.

All three of the auction formats in this case are strategically equivalent to second-price auctions, so the affiliation in values is, in a sense, a red herring – for all three auctions it is still dominant strategy to bid one's value. The two clock auctions are obviously strategyproof, though the AC-B auction still provides bidders with 'less' information than the AC auction. The affiliation hence serves only to complicate the auction for bidders who don't appreciate that the dominant strategy is to bid one's value.

3.2.1 Theoretical and empirical benchmarks

Li (2017)'s experiment delivers results supporting the theoretical framework of obvious strategyproofness – even though the AC and SPSB auctions are strategically equivalent, human subjects tend to be more truthful under the AC auction (which is OSP) than under the SPSB auction. Additional empirical results by Breitmoser & Schweighofer-Kodritsch (2022) show that even OSP itself might not be sufficient in capturing the rich complexities of human behavior – human subjects are less truthful under AC-B than they are under AC (though still more truthful than the SPSB), even though both AC and AC-B are OSP. We replicated these empirical observations using LLMs, suggesting the ability of LLMs to mirror human behavior under these settings.

3.2.2 Simulation evidence

Figure 3 summarizes our findings for the APV setting.

The results of the LLM experiments replicate Li (2017) and Breitmoser & Schweighofer-Kodritsch (2022) closely. However, we see little evidence of learning over time. This is one benchmark that we hope to reach in future iterations of this work. In the experiments in Li (2017) and Breitmoser & Schweighofer-Kodritsch (2022), human subjects improve their understanding of the mechanisms by bidding closer to their true value over time. In the results presented in Figure 3, this effect isn't as pronounced. We suspect that clever prompting strategies, leading to a better and more natural way of communicating the history of past rounds, might lead to more human-like behavior exhibiting learning over rounds.

3.3 Winner's curse

The last set of simulations we ran for single-unit auctions was in common value settings.

Setting: In the common value settings explored here, there are n bidders varying from $n = 2, ..., 6$. Bidders draw values of form $v = c + p$. Once again, the common component is drawn uniformly $c \sim U[0, 79]$ and the private shock component is also drawn uniformly $p \sim U[0, 20]$. However, in the common value setting, bidders have identical ex-post valuations for the good. Hence, agents bid based on values $v = c + p \in [0, 99]$ with a trapezoidal distribution, but only obtain c when they win. The auction is ran as a SPSB auction.

3.3.1 Theoretical and empirical benchmarks

The theoretical optimum here is for bidders to bid $\beta_i(v_i) = \mathbb{E}[c \mid v_i \wedge p_i = p^{(1)}]$, i.e., conditioning both on their value and on the event that their received private shock was the *highest* (it is easy to see the second event is equivalent to the event that the bidder won with strictly monotone $\beta(\cdot)$). The main experiment citation here, Kagel and Levin (1986), makes two positive predictions on the basic common value auction: 1) that even experienced bidders fail to condition on the event where their signal is the highest (called 'item valuation considerations' in their text) thereby still falling victim to the winner's curse, and 2) that the winner's curse barely shows up in small auctions (3-4 bidders) but bites in big auctions (6-7 bidders).

3.3.2 Simulation evidence

We find evidence corroborating both of these predictions. In auctions of all sizes, bidders successfully shade by the expected value of the private shock (i.e., by about $\mathbb{E}[p] = 10)$) but fail to realize that if they won, it is because they drew the *highest* private shock, $p^{(1)}$. As n increases, $\mathbb{E}[p^{(1)}]$ increases, so bidders suffer more in larger auctions. This is demonstrated in Figure 5.

Our evidence suggests that LLMs play at about the level of experienced bidders, generally agreeing quite strongly with existing experimental results. However, 'learning' remains a puzzle in this setting as well – even when told the past history of play, agents don't learn to condition on the event that they win. A defense of LLMs here may be that this is just very hard: learning is non-trivial with human agents as well (in Kagel & Levin (1986) it takes players 15-20 periods to learn). This suggests more sophisticated learning techniques may be required to fully mitigate the winner's curse with LLM agents. It remains to see whether future generations of language models (e.g., GPT5, with better in-context updating and memory) fare better on this front, thereby shading optimally against the winner's curse.

4 Discussion

While this paper focuses on auction theory, future work may use LLM sandboxes to test other kinds of economic mechanisms (e.g., voting, matching, contracts, etc.). As techniques are developed to validate LLM models as proxies for human behavior, they can be used to obtain what would otherwise be prohibitively expensive evidence. As a provocative example, while ethical and financial constraints make it impossible to run voting experiments at the scale of nations, it may be possible to run such experiments with LLM agents.

This paper acts as a proof of concept for LLMs as human proxy agents, with the primary motivation of eventually using LLM agents to inform novel economic design. Some auction formats, such as combinatorial auctions, are complex and can be particularly difficult to run frequently and at scale in traditional laboratory settings. Moreover, with many open questions and limited resources, it can be difficult to triage over the many possible experiments that could be run. Augmenting these traditional lab experiments with LLM experiments, when correctly validated, may open up new avenues in understanding the design tradeoffs in these kinds of complex and often high-stakes environments.

5 Conclusion

This paper reports the results of more than 1,000 auction experiments with LLM agents. In particular, we find behavior that conforms with important experimental results (i.e., evidence of risk-averse bidding, evidence that clock auctions are 'easier' to play, and evidence for the winner's curse in common value settings). We also test theoretical intuitions for combinatorial design in a novel way, generating experimental evidence for three classic CA formats. Though the results are encouraging, we see this work as preliminary, primarily putting forward a framework on how to think about LLM experimental agents as proxy for human agents. In particular, the design space for prompting is large, and we hope that interested readers will use our code to run simulations testing their own prompt variations.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

Gati V Aher, Rosa I Arriaga, and Adam Tauman Kalai. Using large language models to simulate multiple humans and replicate human subject studies. In *International Conference on Machine Learning*, pp. 337–371. PMLR, 2023.

James Brand, Ayelet Israeli, and Donald Ngwe. Using gpt for market research. *Available at SSRN 4395751*, 2023.

Yves Breitmoser and Sebastian Schweighofer-Kodritsch. Obviousness around the clock. *Experimental Economics*, 25(2):483–513, 2022.

Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.

Jiangjie Chen, Siyu Yuan, Rong Ye, Bodhisattwa Prasad Majumder, and Kyle Richardson. Put your money where your mouth is: Evaluating strategic planning and execution of llm agents in an auction arena. *arXiv preprint arXiv:2310.05746*, 2023.

James C Cox, Vernon L Smith, and James M Walker. Theory and individual behavior of first-price auctions. *Journal of Risk and uncertainty*, 1:61–99, 1988.

Sara Fish, Yannai A Gonczarowski, and Ran I Shorrer. Algorithmic collusion by large language models. *arXiv preprint arXiv:2404.00806*, 2024.

John J Horton. Large language models as simulated economic agents: What can we learn from homo silicus? Technical report, National Bureau of Economic Research, 2023.

John H Kagel and Dan Levin. The winner's curse and public information in common value auctions. *The American economic review*, pp. 894–920, 1986.

John H Kagel, Ronald M Harstad, Dan Levin, et al. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica*, 55(6):1275–1304, 1987.

Shengwu Li. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–3287, 2017.

Benjamin S Manning, Kehang Zhu, and John J Horton. Automated social science: Language models as scientist and subjects. Technical report, National Bureau of Economic Research, 2024.

Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pp. 1–22, 2023.

Yixin Wan, George Pu, Jiao Sun, Aparna Garimella, Kai-Wei Chang, and Nanyun Peng. " kelly is a warm person, joseph is a role model": Gender biases in llm-generated reference letters. *arXiv preprint arXiv:2310.09219*, 2023.

A Supplemental material

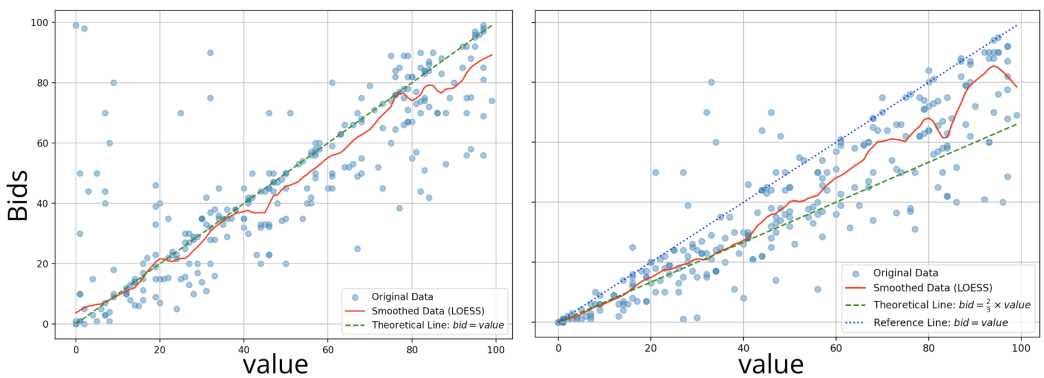

Figure 1: **Comparison of FPSB and SPSB under IPV setting.** The theoretically predicted bid, given that bidders' values are independently drawn from a uniform distribution of [0, 99], is marked in red. The experimental data points are represented by grey triangles. The 45-degree line indicates the scenario where the LLM agents' values equal their bids. The dashed black line represents the LOESS-smoothed data. Left: In FPSB, the experimental bids are ramping up compared to the Bayes Nash prediction. Right: In SPSB, the experimental bids are shading down compared to the dominant strategy.

Comparison	t-statistic	p-value
AC v.s. AC-B	-4.125	6.737e-05
AC v.s. 2P	-5.413	2.101e-07
AC-B v.s. 2P	-5.006	1.043e-06

Table 1: Two-sample t-tests of the mean absolute deviations from bids to values showed significant differences across all comparisons of the strategyproof mechanisms. AC exhibited significantly smaller deviations compared to both AC-B ($t = -4.125$, $p < 0.001$) and 2P ($t = -5.413$, $p < 0.001$). Additionally, the AC-B mechanism showed significantly smaller deviations than the 2P auction ($t = -5.006$, $p < 0.001$).

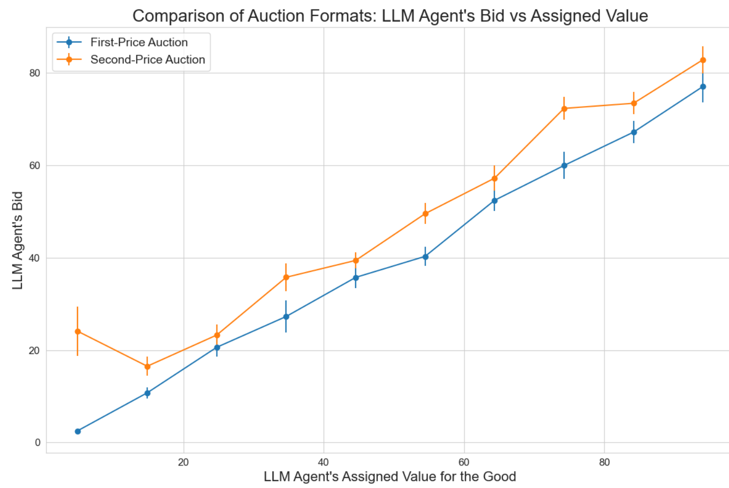

Figure 2: **Cross-format Comparison of FPSB and SPSB under IPV setting.** The bins are organized in $10 increment. The orange curve stands for the Second-Price Sealed-Bid auction and the blue curve stands for the First-Price Sealed-Bid auction. The x-axis represents the assigned value of the good to the LLM agent, ranging from low to high values. The y-axis shows the LLM agent's bid amount. The blue line represents First-Price auctions while the orange one represents Second-Price auctions. The error bars indicate the standard error of the mean bid at each value point, showing the variability in bidding behavior.

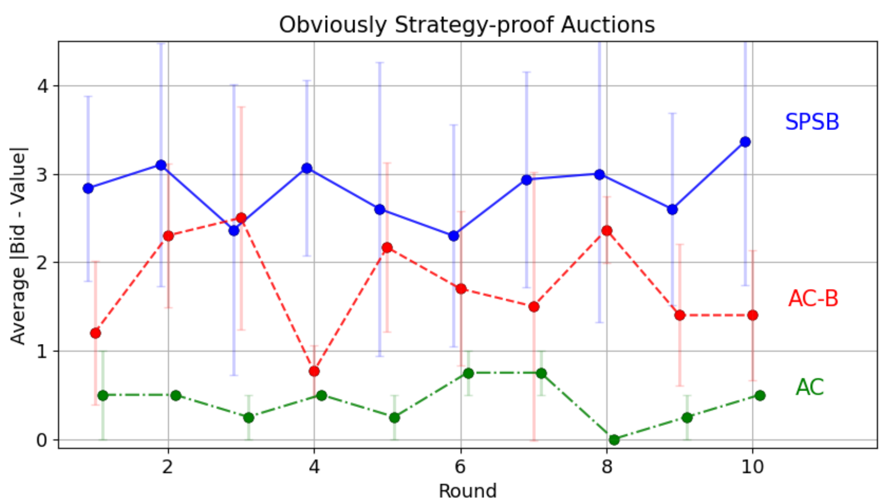

Figure 3: **Comparison of three strategically equivalent auctions.** Ascending-clock (AC) and its variant without dropping-out information (AC-B) are obviously strategy-proof while second-price sealed-bid (SPSB) is not. Here, the green dot-dash line plotted the mean absolute deviation from bids to values in AC. Red dash line is for AC-B. And blue solid line is for SPSB. The mean absolute deviations are smallest in AC and highest in SPSB and with AC-B in between, consistent with the human lab experiments in Breitmoser & Schweighofer-Kodritsch (2022). However, there is no sign of learning over rounds.

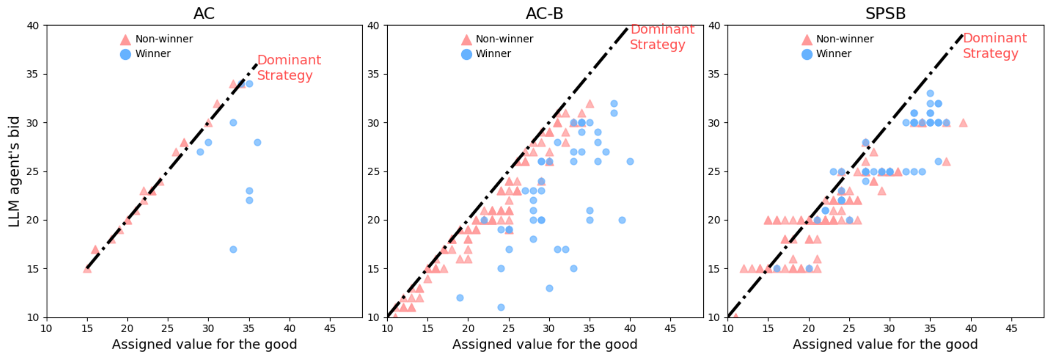

Figure 4: **Bid versus value scatter plot for different auction formats.** The X-axis is the assigned value for the good and Y-axis is LLM agent's bid. Red triangle stands for the bids of non-winner in the auction while the blue circle represents the ones of the winners. In the Ascending Clock (AC) auction, all players dropped out at the dominant strategy price within one round of bidding. The bids in AC-Blind auction are close to the optimal line, but showed a significant pattern of early dropouts. In the Second-Price Sealed-Bid auction, almost all the bids deviated from the dominant strategy, going both higher and lower than the optimal price.

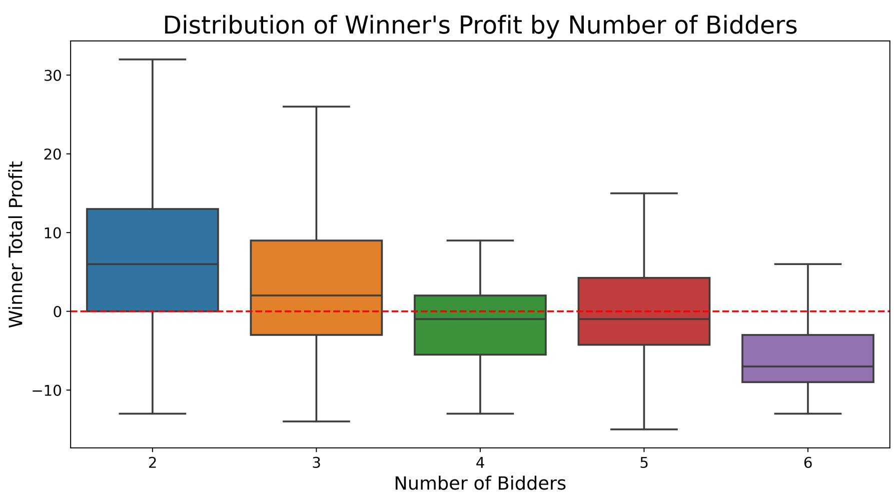

Figure 5: **Distribution of the winner's total profit across auctions with 2 to 6 bidders.** Each box shows the interquartile range of profits, with the median indicated by the central line. The horizontal red dashed line represents zero profit. As the number of bidders increases, the median winner's total profit decreases and more frequently turns negative. This echoes with the intensifying effect of the winner's curse in larger auctions.

B Prompts in IPV section

First-Price Sealed-Bid (FPSB) with IPV

In this game, you will participate in an auction for a prize against num]_bidders other bidders. At the start of each round, bidders will see their value for the prize, randomly drawn between $0and$private, with all values equally likely.
After learning your value, you will submit a bid privately at the same time as the other bidders. Bids must be between \$0 and \$private in \$increment increments. The highest bidder wins the prize and pays their bid amount. This means that, if you win, we will add to your earnings the value for the prize, and subtract from your earnings your bid. If you don't win, your earnings remain unchanged.
After each auction, we will display all bids and profits. Ties for the highest bid will be resolved randomly.

Second-Price Sealed-Bid (SPSB) with IPV

In this game, you will participate in an auction for a prize against num_bidders other bidders. At the start of each round, bidders will see their value for the prize, randomly drawn between $0and$private, with all values equally likely.
After learning your value, you will submit a bid privately at the same time as the other bidders. Bids must be between \$0 and \$private in \$increment increments. The highest bidder wins the prize and pays the second-highest bid. This means that, if you win, we will add to your earnings the value for the prize, and subtract from your earnings your bid. If you don't win, your earnings remain unchanged.
After each auction, we will display all bids and the winner's profits. Ties for the highest bid will be resolved randomly.

C Prompts in OSP section

Second-Price Sealed-Bid (SPSB) with APV

In this game, you will bid in an auction for a prize against num_bidders other bidders. The prize may have a different dollar value for each person in your group. You will play this game for n rounds. All dollar amounts in this game are in increment increments. At the start of each round, we display your value for this round's prize. If you win the prize, you will earn the value of the prize, minus any payments from the auction.
Your value for the prize will be calculated as follows: 1. For each group we will draw a common value, which will be between common_low and common_high. Every number between common_low and common_high is equally likely to be drawn. 2. For each person, we will also draw a private adjustment, which will be between 0 and private. Every number between 0 and private is equally likely to be drawn. In each round, your value for the prize is equal to the common value plus your private adjustment. At the start of each round, you will learn your total value for the prize, but not the common value or the private adjustment. This means that each person in your group may have a different value for the prize. However, when you have a high value, it is more likely that other people in your group have a high value.
The auction proceeds as follows: First, you will learn your value for the prize. Then you can choose a bid in the auction. Each person in your group will submit their bids privately and at the same time. All bids must be between min_price and max_price, and in increment USD increments. The highest bidder will win the prize, and make a payment equal to the second-highest bid. This means that we will add to her earnings her value for the prize, and subtract from her earnings the second-highest bid. All other bidders' earnings will not change.

Blind Ascending-Clock (AC-B) with APV

In this game, you will bid in an auction for a prize against num_bidders other bidders. The prize may have a different dollar value for each person in your group. You will play this game for n rounds. All dollar amounts in this game are in increment increments. At the start of each round, we display your value for this round's prize. If you win the prize, you will earn the value of the prize, minus any payments from the auction.
Your value for the prize will be calculated as follows: 1. For each group we will draw a common value, which will be between common_low and common_high. Every number between common_low and common_high is equally likely to be drawn. 2. For each person, we will also draw a private adjustment, which will be between 0 and private. Every number between 0 and private is equally likely to be drawn. In each round, your value for the prize is equal to the common value plus your private adjustment. At the start of each round, you will learn your total value for the prize, but not the common value or the private adjustment. This means that each person in your group may have a different value for the prize. However, when you have a high value, it is more likely that other people in your group have a high value.
The auction proceeds as follows: First, you will learn your value for the prize. Then, the auction will start. We will display a price to everyone in your group, that starts low and counts upwards in increment USD increments, up to a maximum of max_price. At any point, you can choose to leave the auction. The starting bidding will be min_price. When there is only one bidder left in the auction, that bidder will win the prize at the current price. This means that we will add to her earnings her value for the prize, and subtract from her earnings the current price. All other bidders' earnings will not change. At the end of each auction, we will show you the prices where bidders stopped, and the winning bidder's profits. If there is a tie for the highest bidder, no bidder will win the object.

Open Ascending-Clock (AC) with APV

In this game, you will bid in an auction for a prize against num_bidders other bidders. The prize may have a different dollar value for each person in your group. You will play this game for n rounds. All dollar amounts in this game are in increment increments. At the start of each round, we display your value for this round's prize. If you win the prize, you will earn the value of the prize, minus any payments from the auction.
Your value for the prize will be calculated as follows: 1. For each group we will draw a common value, which will be between common_low and common_high. Every number between common_low and common_high is equally likely to be drawn. 2. For each person, we will also draw a private adjustment, which will be between 0 and private. Every number between 0 and private is equally likely to be drawn. In each round, your value for the prize is equal to the common value plus your private adjustment. At the start of each round, you will learn your total value for the prize, but not the common value or the private adjustment. This means that each person in your group may have a different value for the prize. However, when you have a high value, it is more likely that other people in your group have a high value.
The auction proceeds as follows: First, you will learn your value for the prize. Then, the auction will start. We will display a price to everyone in your group, that starts low and counts upwards in increment USD increments, up to a maximum of max_price. At any point, you can choose to leave the auction. The starting bidding will be min_price. When there is only one bidder left in the auction, that bidder will win the prize at the current price. This means that we will add to her earnings her value for the prize, and subtract from her earnings the current price. All other bidders' earnings will not change. At the end of each auction, we will show you the prices where bidders stopped, and the winning bidder's profits. If there is a tie for the highest bidder, no bidder will win the object.

D Prompts in Common value auction

Second-Price Sealed Bid with common value

In this game, you will bid in an auction for a prize against num_bidders other bidders. The prize may have a different dollar value for each person in your group. You will play this game for n rounds. All dollar amounts in this game are in increment increments. At the start of each round, we display your value for this round's prize. If you win the prize, you will earn the value of the prize, minus any payments from the auction.
Your value for the prize will be calculated as follows: 1. For each group we will draw a common value, which will be between common_low and common_high. Every number between common_low and common_high is equally likely to be drawn. 2. For each person, we will also draw a private adjustment, which will be between 0 and private. Every number between 0 and private is equally likely to be drawn. In each round, your value for the prize is equal to the common value. However, at the start of each round, you will learn only your total value for the prize, not the common value or the private adjustment. This means that each person in your group may have a different perceived value for the prize.
The auction proceeds as follows: First, you will learn your (perceived) value for the prize. Then you can choose a bid in the auction. Each person in your group will submit their bids privately and at the same time. All bids must be between min_price and max_price, and in increment USD increments. The highest bidder will win the prize, and make a payment equal to the second-highest bid. This means that we will add to her earnings the (true) common value for the prize, and subtract from her earnings the second-highest bid. All other bidders' earnings will not change.