Learning from Synthetic Labs: Language Models as Auction Participants *

Anand Shah* MIT
Kehang Zhu* Harvard
Yanchen Jiang Harvard
Jeffrey G. Wang Harvard

Arif K. Dayi Harvard
John J. Horton MIT & NBER
David C. Parkes Harvard

March 21, 2025

Abstract

This paper investigates the behavior of simulated AI agents (large language models, or LLMs) in auctions, introducing a novel synthetic data-generating process to help facilitate the study and design of auctions. We find that LLMs - when endowed with chain of thought reasoning capacity - agree with the experimental literature in auctions across a variety of classic auction formats. In particular, we find that LLM bidders produce results consistent with risk-averse human bidders; that they perform closer to theoretical predictions in obviously strategy-proof auctions; and, that they succumb to the winner's curse in common value settings. On prompting, we find that LLMs are not very sensitive to naive changes in prompts (e.g., language, currency) but can learn to improve dramatically towards theoretical predictions with the right mental model (i.e., the language of Nash deviations). We run 1,000+ auctions for less than $400 with GPT-4 models (three orders of magnitude cheaper than modern auction experiments) and develop a framework flexible enough to run auction experiments with any LLM model and a wide range of auction design specifications, facilitating further experimental study by decreasing costs and serving as a proof-of-concept for the use of LLM proxies.

*This research was made possible by a generous grant from Dropbox Inc. Thanks to Aidan Toner-Rodgers and Parker Whitfill for their helpful feedback. Author's contact information, code, and data will be available at https://github.com/KeHang-Zhu/llm-auction. Both Anand V. Shah and Kehang Zhu contributed equally to this work. John J. Horton is a co-founder of Expected Parrot Inc., which develops generative AI models for market research applications. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

1 Introduction

Rapidly advancing capabilities of large language models (LLMs) raise an exciting possibility: might LLM agents serve as cost-effective proxies for human bidders, thereby making experiments orders of magnitude cheaper to run? In this paper, we contribute to an emerging literature on LLM proxies by testing whether LLM agents can replicate key empirical regularities of the auctions literature. We map LLM behavior in standard auction formats against established human and theoretical benchmarks from economics, providing a proof-of-concept for low-cost, large-scale auction experimentation using LLM bidders.

We begin by benchmarking play with LLM bidders in classic environments: sealed-bid auctions in independent, private-value settings. When there are departures from the theory, we are interested in whether the departures agree with empirical results (Kagel and Levin, 1993; Kagel and Roth, 2020). We find that agents fail to replicate revenue equivalence results (as existing empirical evidence suggests). We find that bids in the SPSB auction are higher than bids under the FPSB auction, as would be expected, but that there is a smaller separation between the two than predicted by theory. This is primarily due to bids under the FPSB auction being higher than the risk-neutral Bayes-Nash equilibrium suggests. One possible explanation is that LLMs play according to some level of risk aversion: this is consistent with Cox et al. (1988)'s survey of over 1,500 IPV auction experiments, which finds that revenue under the FPSB auction is higher than under the SPSB auction and suggests that this is due to risk attitudes. In addition, we see LLMs in the SPSB auction tend to submit bids lower than their value to an extent that is not echoed in the experimental economics literature. We also supplement this analysis with simulations of the third-price sealed bid (TPSB) and all-pay auctions.

Next, we fix the second-price auction and consider classic ways to make the auction 'easier' or 'harder' to play, to learn whether LLMs exhibit behavior similar to humans in the face of cognitive constraints.

A classic puzzle in the study of second price auctions is that bidders find it easier to play the dominant strategy equilibrium in second-price ascending clock auction rather than second-price sealed-bid auctions, despite the two having the same strategic form (Kagel et al., 1987). Theory and empirics (Li, 2017) have recently been developed to argue this is because ascending clock formats are less cognitively demanding than sealed-bid ones, causing humans to make less mistakes relative to the dominant strategy of bidding one's value. We find that LLMs also find clock formats easier to play – they are more likely to exit the auction (the analogue to bidding in the sealed-bid format) at their value. Following Li (2017), we study this in an affiliated private value setting, replicating their results with LLMs. We also compare the typical clock auction with a blind variant (a clock auction where bidders do not know when others drop out) and the corresponding second-price sealed-bid auction, and find that switching to a clock format accounts for most of the improvement in LLM play.

Next, we add complexity to the auction environment by considering the second-price auction under a common value setting. Under common values, all agents agree on their value for the prize c, but each only observes their valuation for the prize under some idiosyncratic, private measurement error $c + p_i$. This auction environment is more cognitively demanding than the independent, private values case as agents do not know whether their observed high valuations are due to high draws of c or the measurement error p_i, and this tension produces the winner's curse: that the winning agents wish they had never won, having been adversely selected to be the bidder who least appreciated their high signal was due to a high p_i. The winner's curse is purely a non-rational phenomenon – only the cognitively limited agent suffers from it. The existing economics literature finds humans regularly suffer the winner's curse in common value settings, and we find results vindicating this evidence with LLM agents: LLMs succumb to the winner's curse as well. In particular, Kagel and Levin (1986) famously find experimental evidence for the winner's curse and argue that the winner's curse barely bites with 3-4 bidders but seriously bites for larger auctions with 6+ bidders.

We find qualitatively similar results hold with LLM agents as well.

These classic results are some of the reasons economists have found auction theory to be beautiful over the last few decades, but auctions can also be useful (Milgrom and Segal, 2020). Auctions are used to clear billions of dollars of value annually, and one may ask how LLMs fare in less controlled settings. To answer this, we design a more realistic auction environment inspired by eBay's online marketplace. We examine two design parameters: a hidden reserve price and a modified closing rule (soft-close, which extends the auction when there is bidding activity). First, we demonstrate that LLM agents replicate emergent, real-world bidding behaviors seen on eBay with no prompting: in auctions with a hard closing time, LLMs frequently delay bids until the last second to 'snipe' their competitors as humans would. Second, experiments testing the design of an extended closing rule greatly reduce bid sniping and improve price discovery, echoing dynamics of online auctions between Amazon and eBay in the early 2000s (Roth and Ockenfels, 2002).

The ability to interface more directly with our experimental subjects also enables us to study the explanation of mechanisms more carefully. Auctions are most prized as mechanisms to allocate goods efficiently, but if agents don't understand the rules they can't be expected to play well. How can we describe and explain mechanisms to empower optimal play cheaply? We conclude the body of our paper by testing interventions designed to influence or aid LLM bidders' understanding of the underlying economics of the setting and its willingness to follow rules. In particular, we test six interventions: an extensive-form menu description meant to simplify the equilibrium logic inspired by Gonczarowski et al. (2023); a description describing the auction via a clock (without actually running the clock auction) inspired by Breitmoser and Schweighofer-Kodritsch (2022); a description walking an LLM through the logic of Nash deviations; a description which prompts agents to be risk-neutral; a description which reveals the correct statement that the SPSB's dominant strategy is to bid one's value; and a description that incorrectly reveals a false statement that the SPSB's dominant strategy is to bid half of one's value. We find that all the interventions that were meant to improve the correctness of an LLM agent's bidding strategy helped, with the logic of Nash deviations helping the most of the first three. Unsurprisingly, the false statement intervention causes LLMs to play poorly, but interestingly, the LLM continues to experiment with new strategies upon determining the given advice was bad.

The remainder of the paper proceeds as follows. In Section 1.1, we report related work. In Section 2, we describe our simulation procedure and prompting methods in more depth. In Section 3.1, we benchmark LLM play against sealed-bid auctions in independent, private value settings. Then, for the rest of the paper, we fix the game as a second-price auction and test various design considerations. In Section 3.2, we vary the auction format and find that LLMs play clock formats the best. In Section 3.3, we add cognitive complexity to the auction environment by switching to a common value environment and show that LLMs succumb to the winner's curse. Both of these Sections replicate key empirical regularities found in laboratory and real-life play. In Section 4, we implement auctions inspired by eBay's marketplace and provide new evidence for classic debates on design parameters in internet auctions. In Section 5, we introduce interventions designed to influence or aid LLM bidders' understanding of the underlying economics of the setting. We find that prompting matters, and that, surprisingly, LLMs improve play dramatically when instructed to follow the logic of Nash deviations.

In the Appendix we report complete simulation procedures, prompts and numerous robustness checks. In particular, we ran experiments with prompts from other contemporaneous auctions experiments (e.g., the Appendix script from Li (2017)), with prompts in different languages and with different currencies, and with varying numbers of bidders. All experiments are run with chain-of-thought agents, but we also run an ablation study with out-of-the-box agents. We also find interesting evidence that 'goal' prompting (e.g., reminding the LLM that the goal is to maximize profit) leads to bidding that more closely follows rational economic theory, leading to higher allocative efficiency. See also Manning et al. (2024) for a related

discussion.

To obtain the data for these empirical results, we have developed a code repository to systematically run experiments with a certain number of bidders and a particular prompt. In particular, our repository is flexible enough that it can be used to generate synthetic data for almost any describable auction format and including auctions with single or multiple goods.[1] For the experiments described in this paper, we ran more than $1,000$ auctions with more than $5,000$ GPT-4 agent participants for costs totaling less than $\$400$. In contrast, the largest survey of auction experiments to date comes from Cox et al. (1988), reporting the results of running $1,500+$ auctions and with considerably higher total cost.

1.1 Related Work

Auctions: There's a vast quantity of theoretical and experimental literature on auctions. While Krishna (2009)'s textbook and Kagel and Roth (2020)'s handbook provided invaluable general resources, we will stay focused only on the citations relevant to the results in this paper.

Starting with the IPV case, the benchmark of revenue equivalence in the risk-neutral case is exposited in the seminal Myerson (1981). The experimental evidence for departures due to risk-aversion in the FPSB auction is thoroughly documented in Coppinger et al. (1980) and Cox et al. (1988)'s survey. The experimental evidence for the common error of bidding above one's value in the SPSB auction is documented in Kagel and Levin (1993).

Recent work on obvious strategy-proofness began with Li (2017), who demonstrates empirically that human subjects tend to be less truthful in second-price sealed-bid auctions than ascending clock auctions in the APV setting, even though the two auctions are strategically equivalent. Li also provides a theoretical framework for the results. To better understand our simulations, we also consider the experimental evidence presented by Breitmoser and Schweighofer-Kodritsch (2022), who investigate intermediate auction formats that decompose the behavioral effects in Li (2017).

For common value settings, the winner's curse has been documented empirically and experimentally. Experimental evidence for the winner's curse was first given by Bazerman and Samuelson (1983). For the present paper, we primarily follow the empirical approach in Kagel and Levin (1986) (FPSB auctions with varying numbers of bidders) and Kagel et al. (1987) (SPSB auctions with informational interventions). Empirically, the winner's curse was first documented in 1971, in auctions for oil drilling rights by Capen et al. (1971). Theoretically, while the winner's curse is fundamentally a problem of non-rational play, Charness and Levin (2009) provide a model to cast the adverse-selection problem as a failure from cognitive constraints, microfounding our interest in evaluating LLM susceptibility to the winner's curse via bounded rationality.

LLMs as simulated agents: Recent LLMs, having been trained on an enormous corpus of human-generated data, are able to generate human-like text and reasoning (Achiam et al., 2023; Bubeck et al., 2023). Yet, they are far from perfect – in particular, displaying limited planning abilities and reflecting various cognitive biases endemic to human agents (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; Brand et al., 2023; Park 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. Manning et al. (2024) enable multiple GPT-4 agents to interact and simulate various social science scenarios, including bargaining, job interviews, and bail-setting. Finally, Raman et al. (2024) develop a benchmark of rational play for LLMs over a broad

[1] We will make the code-base public and hope this will facilitate additional research into the role of LLMs as modeling human bidder behavior.

range of economic tasks. Compared to humans, the inference cost of LLM queries is very low and continues to decrease (Achiam et al., 2023; Bae et al., 2023; Patel et al., 2023).

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 to study dynamic decision making. Finally, Manning et al. (2024) run a more limited study on a variant of the ascending clock auction with three bidders, focusing on deviations from rational economic theory in considering bidders' values and the final clearing price.

2 Methods

2.1 LLM as Auction Participants

We use a single LLM API call to represent one bidder in our auction simulations. Unlike traditional agent-based modeling, we do not anchor bidder behaviors or personas mechanically beyond a name to index them (e.g., Alice, Betty or Charles). Instead, we allow LLMs to conduct their own strategic reasoning and in-context learning to generate bids. Hence, we focus on how these LLM agents strategically interact, without concern for demographic representativeness, as our primary interest is strategic behavior in mechanisms.

To ensure that LLM agents make profit-maximizing decisions, we append a prompt prefix that explicitly instructs them to seek long-run profit maximization. This instruction is universal in all the experiments ran in this paper, and is inspired by Fish et al. (2024). Complete details on prompts and execution can be found in Appendix B.1. Specific implementation details for non-sealed-bid formats (the clock and eBay auction formats) are provided in their respective sections.

Before eliciting bids in each period, we provide a description of the auction rules along with randomly drawn values and information from previous rounds. Standard prompting techniques like chain-of-thought (Wei et al., 2022) are also used to mimic the strategic 'thinking period' which is common to auction experiments in the literature (e.g., Li (2017)).

2.2 Simulation Procedure

For each auction experiment, we simulate n (often 3) LLM agents bidding against one another. Each setting is repeated multiple times, with randomized private values drawn for each round. We design our multi-round protocol to mirror standard laboratory experiments in the auction literature. The item to be bid in standard auctions is specified as a "prize". In eBay auctions, we arbitrarily set the prize item as a "256GB IPhone 16 pro" and the item condition is set to be "used" for realism.

Below is our simulation procedure for each auction; full prompt texts are in Appendix G, and an example round is in Appendix B.1. The simulation procedure is designed to mimic experiments in the literature as closely as possible:

1. *Planning.* LLM agents receive a description of the auction format and are asked to propose a bidding plan.
2. *Value Realization.* Each LLM agent is assigned a random private value.
3. *Bidding.* Given their value, each agent places a bid according to their stated plan.
4. *Outcome.* The system computes allocations and payments. Each LLM agent learns its own profit or loss outcome, as well as the full set of bids submitted (but not other agents' private values).

5. *Reflection.* Before the next round, each LLM agent reviews the results and possibly revises its approach. These updates are stored in the `HISTORY` variable.

6. *Repetition.* One complete "plan-bid-reflect" loop is one round. We repeat 15 rounds in each experiment, with the entire `HISTORY` carried forward.

API calls are made using the Python package Expected Parrot Domain-Specific Language (EDSL), a wrapper that structures LLM queries and responses in the desired format (Horton et al., 2024). Each agent in our sample is supported by a separate API call via EDSL, thereby nullifying the risk of collusion. For every auction, we use models from OpenAI at a temperature of 0.5.[2] The full code base will be made open-source for research and replication at GitHub upon publication.

3 Benchmarking Against Existing Auctions

3.1 Sealed Bid Auctions with Independent Private Values

We begin by examining the First-Price Sealed-Bid (FPSB), Second-Price Sealed-Bid (SPSB), Third-Price Sealed-Bid (TPSB) and All-Pay Sealed-Bid auctions in an independent private values (IPV) setting.

3.1.1 Setting

There are three bidders in each auction, and each bidder i's value is drawn from an independent, uniform distribution $v_i \sim U[0, 99]$. Bidders, upon observing their value, submit a sealed bid $\beta(v)$, with $\beta(\cdot)$ a vector mapping each component value to its corresponding bid, corresponding to a strategy profile for reporting based on the valuation profile.

In the FPSB auction, the highest bidder pays her bid and receives the prize (and all other bidders pay 0 and receive no prize). Formally, the payment for an agent is given by:

$$t_i(\beta(v)) = \mathbb{1}_{i \text{ won the auction}} \cdot \beta_i(v).$$

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). The payment for an agent is given by:

$$t_i(\beta(v)) = \mathbb{1}_{i \text{ won the auction}} \cdot \beta^{(2)}(v),$$

where $\beta^{(2)}$ represents the second-order statistic or the second-largest bid. Analogously, in the TPSB auction, the highest bidder pays the third-highest bid and receives the prize (and all other bidders pay 0 and receive no prize). The payment for an agent is given by:

$$t_i(\beta(v)) = \mathbb{1}_{i \text{ won the auction}} \cdot \beta^{(3)}(v),$$

where $\beta^{(3)}$ represents the third-order statistic or the third-largest bid. Finally, in the all-pay auction, all bidders pay their bid but only the highest bidder wins the prize. The payment for an agent in the all-pay auction is given by:

$$t_i(\beta(v)) = \beta_i(v).$$

Bids are submitted in \$1 increments and ties are resolved randomly.

[2] We set a non-zero temperature to encourage diversity in strategies and reflections. No parameter tuning was performed to achieve these results.

3.1.2 Prompts

Following the simulation procedure outlined in Section 2.2, we employ the following prompts for the four auctions tested in this Section. For each prompt, only the third paragraph (describing the allocation and transfer rule) differ. For the FPSB, we describe the auction format as follows:

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their value for
   the prize, randomly drawn between \$0 and \${{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. If
    you win, your earnings will increase by your value for the
   prize, and decrease by your bid. If you don't win, your
   earnings will remain unchanged.
After each auction, we will display all bids and the winner's
   profits. Ties for the highest bid will be resolved randomly.
```

For the SPSB, we describe the auction format as follows:

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their value for
   the prize, randomly drawn between $0 and ${{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. If you win, your earnings will increase by your value
   for the prize, and decrease by the second-highest bid. If
   you don't win, your earnings will remain unchanged.
After each auction, we will display all bids and the winner's
   profits. Ties for the highest bid will be resolved randomly.
```

For the TPSB, we describe the auction format as follows:

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their value for
   the prize, randomly drawn between $0 and ${{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 third-highest
   bid. If you win, your earnings will increase by your value
   for the prize, and decrease by the third-highest bid. If you
    don't win, your earnings will remain unchanged. If there
   are less than three bids, no one will win the auction.
After each auction, we will display all bids and the winner's
   profits. Ties for the highest bid will be resolved randomly.
```

For the all-pay auction, we describe the auction format as follows:

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their value for
   the prize, randomly drawn between $0 and ${{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. All bidders (including the
   winner) pay their submitted bid. If you win, your earnings
   will increase by your value for the prize, and decrease by
   your bid. If you don't win, your earnings will still
   decrease by your bid.
After each auction, we will display all bids and all bidders'
   profits. Ties for the highest bid will be resolved randomly.
```

3.1.3 Theoretical benchmarks

The risk-neutral equilibrium strategy in all four of these auction formats is well understood. We summarize the functional forms of these strategies - specific to our setting - below in Table 1.

Risk-neutrality is a particularly strong theoretical assumption. In practice, human subjects deviate from the risk-neutral equilibria of Table 1, and these deviations have motivated a broader theoretical literature seeking to rationalize bidding behavior under alternative assumptions about bidder preferences and decision-making (ambiguity aversion, loss aversion, probability misperception, etc). The most enduring and broadly applicable explanation remains risk aversion – we will focus on this explanation in the following Section as well.

Risk-averse bidders are willing to lose value in expectation for more certainty. In this paper, we focus on two canonical formalizations: Constant Absolute Risk Aversion (CARA) and Constant Relative Risk

Auction Type	Risk-Neutral Equilibrium Strategy
First-Price Sealed-Bid*	$\frac{2}{3}v$
Second-Price Sealed-Bid**	v
Third-Price Sealed-Bid*	$2v$
All-Pay Auction*	$\frac{2}{3 \cdot 99^2}v^3$

Table 1: Risk-neutral equilibria of IPV auctions with 3 bidders and values drawn uniformly, $v \sim U[0, 99]$. A single asterisk (*) denotes a Bayes-Nash equilibrium, while a double asterisk (**) denotes a dominant strategy equilibrium. Derivations in Appendix A.2.

Aversion (CRRA) preferences.[3]

$$\text{CARA:} \quad U(x) = -e^{-\lambda x}, \quad \lambda > 0$$
$$\text{CRRA:} \quad U(x) = \frac{x^{1-\rho}}{1-\rho}, \quad \rho > 0, \quad \rho \neq 1$$

The CARA specification assumes a bidder's absolute risk tolerance remains fixed across wealth levels, whereas CRRA assumes risk tolerance scales with wealth. CARA is more analytically tractable (and has a closed form for each auction) but is not scale invariant (switching units from dollars to cents changes CARA parameter estimates). CRRA is less analytically tractable (we report numerical solutions for empirical results) and scale invariant, making it easier to compare our results with previous empirical results.

Table 2 presents the equilibrium bidding strategies under CARA and CRRA preferences, with functional forms that explicitly incorporate the risk-aversion parameter. There is no new theoretical prediction for the SPSB auction under risk-aversion, as dominant strategies are invariant to risk attitudes.

Auction Type	CARA Equilibrium	CRRA Equilibrium
First-Price Sealed-Bid	$v - \frac{1}{\lambda}\ln(3 - 2e^{-\lambda v})$	$\frac{2}{2+\rho}v$
Third-Price Sealed-Bid	$v + \frac{1}{\lambda}\ln(1 + \lambda v)$	Numerical solution to ODE
All-Pay Auction	Numerical solution to ODE	Numerical solution to ODE

Table 2: Bayes-Nash equilibria under CARA (λ) and CRRA (ρ) preferences with n bidders and values drawn uniformly, $v \sim U[0, 1]$. Derivations in Appendix A.3.

Finally, irrespective of risk attitudes, theory universally predicts monotone bidding functions in single-unit IPV settings. Monotone bidding is a benchmark of auction theory more generally, reflecting the assumption in utility theory that "more is better".

[3]The CARA specification is also given as $U(x) = 1 - e^{-\lambda x}$, the same preferences under affine transformation.

3.1.4 Empirical benchmarks and results

For each of the IPV auction formats, we perform 5 experiments with 15 rounds and 3 LLM bidders competing over a prize.[4]

Our main results are plotted in Figure 1 and Figure 2, alongside the prediction of rational theory. On each plot, we fit a Loess-smoothed curve – these curves impose no parametric restrictions on the data and constitute evidence of monotonicity. A formal test is also offered later this section. We will first discuss the results of each of our simulations compared to the existing literature, highlighting five points in the analysis. 1) LLM play, in aggregate, looks similar to human play. 2) However, individual bids from LLMs are often mistakes, far from the optimal play. 3) These mistakes are consistent with risk aversion. This argument is made precise by identifying CARA and CRRA utility coefficients for LLM bidding and comparing them to the literature. 4) Mistakes in the SPSB, though unattributable to risk attitudes, also occur at the same rate as in human literature with some qualitative differences. 5) In sum, we find the largest qualitative difference between human and LLM bidding is that LLM bidding is higher variance. This motivates some interest in prompting which we return to later in the paper.

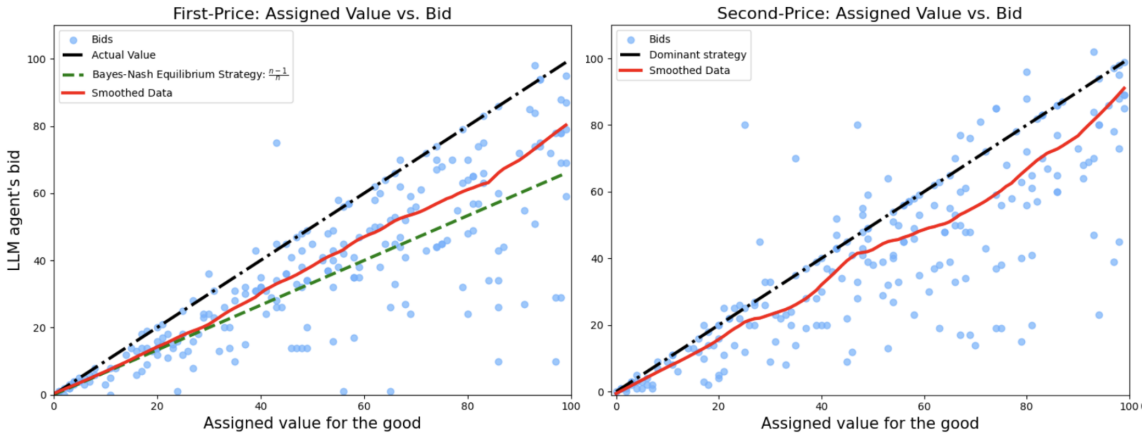


Figure 1: **Comparison of bids vs. values for the FPSB (left) and SPSB (right) auction in an IPV setting.** Bidders' values are independently drawn from a uniform distribution over [0, 99] and LLM bids, given a value, are represented by blue circles. The black dotted line represents the 45° line (the dominant strategy for SPSB auctions), the green dashed line represents the Bayes-Nash Equilibrium on the FPSB plot, and the solid red line represents the LOESS-smoothed bidding curve for both plots.

First, we summarize the distributional features of our sealed-bid auctions. Bidding is diffuse, with most bidders deviating from optimal play. Table 3 summarizes these distributions, categorizing bids as underbids, value bids, or overbids relative to one's value. However, aggregate patterns across auction formats provide encouraging evidence that, in aggregate, LLM agents are responding strategically to the change in auction format in ways theory predicts. First, we find evidence of monotonic bidding (in aggregate), testing this with a Kendall-τ test in Table 5. Monotonicity is an empirical regularity in the auctions literature, and the

[4]While these specifications are designed to meet the 'classic' auction formats of the literature, we also report the results of several variants in Appendix D. In particular, we run one-shot auction simulations (as opposed to multi-round) with $n = 4, 5$, and 10 bidders, auctions with prompts using various currencies, and auctions with non-English prompts using various currencies.

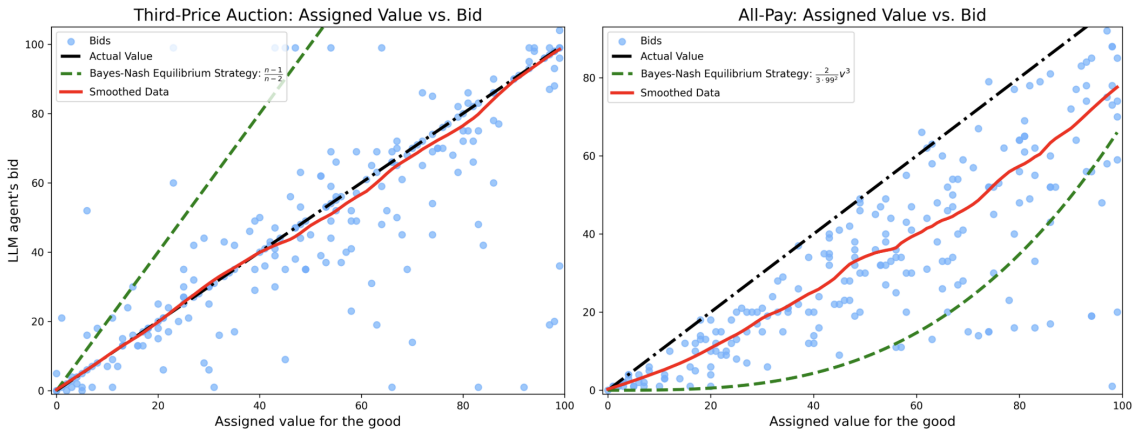


Figure 2: **Comparison of TPSB and all-pay under IPV setting.** Bidders' values are independently drawn from a uniform distribution of [0, 99] and the experimental data points are represented by blue circles. The black dotted line represents bid=value (the dominant strategy for SPSB auctions), the green line on the FPSB figure denotes the Bayes-Nash Equilibrium, and the solid red line represents the LOESS-smoothed bidding data. Here, we show results for third-price auctions on the left, and all-pay auctions on the right.

basis for the ubiquitous monotonicity assumption taken in auction theory. Second, we find evidence that as k increases for the k-th price auction, the distribution of bids versus values shifts upward (Monderer and Tennenholtz, 2000). In the FPSB auction, only 4.4% of bids exceed the bidder's value (a weakly dominated strategy), with the vast majority of bidders underbidding relative to their values. In contrast, the distribution shifts upward for SPSB (12.44% bids higher than their value) and TPSB (27.11% bids higher than their value). χ^2 tests between the bidding distributions confirm this visual intuition, with results reported in Table 4. The closest evidence from the empirical literature, Kagel and Levin (1993), find similar results, also running χ^2 tests to confirm that distributions are distinct and moving relatively up from the FPSB to the SPSB to the TPSB auction. The difference between Kagel and Levin (1993)'s results and our own on this matter is one of magnitudes – while the bidding distributions for the $k-$price auction are significantly distinct and increasing in k, there is more separation between bidding distributions with human bidders than LLM bidders.

Auction Type	**Underbidding (%)**	**Value Bidding (%)**	**Overbidding (%)**
First-Price Sealed-Bid	82.22	13.33	4.44
Second-Price Sealed-Bid	65.78	21.78	12.44
Third-Price Sealed-Bid	42.22	30.67	27.11
All-Pay	94.67	4.44	0.89

Table 3: **Bidding relative to one's value.** Underbidding is defined as $\beta_i(v) < v_i$; Value bidding is defined as $\beta_i(v) = v_i$; Overbidding is defined as $\beta_i(v) > v_i$.

Despite features in aggregate corresponding to rational theory, most individual bidding is not optimal

Auction 1	Auction 2	χ^2 Statistic	p-value
First-Price	Second-Price	22.97	0.0063
First-Price	Third-Price	35.68	< 0.0001
Second-Price	Third-Price	36.67	< 0.0001
All-Pay	First-Price	11.77	0.2642
All-Pay	Second-Price	24.69	0.0033
All-Pay	Third-Price	59.46	< 0.0001

Table 4: **Chi-square tests for sealed bid auctions in an IPV environment.** All auctions with the exception of the all-pay and first-price auctions are statistically distinct.

Auction Type	Kendall's τ_b	p-value
First-Price	0.672	< 0.0001
Second-Price	0.666	< 0.0001
Third-Price	0.671	< 0.0001
All-Pay	0.660	< 0.0001

Table 5: **Kendall's τ_b test for bid monotonicity.** $\tau_b = 1$ indicates perfect positive monotonicity, while $\tau_b = -1$ indicates perfect negative monotonicity.

play. We can account for many of these mistakes by modeling risk preferences. To quantify risk attitudes in LLM bidding, we estimate risk coefficients under CARA and CRRA preferences for the FPSB, TPSB and all-pay auctions according to the specification in Section 3.1.3. The SPSB is excluded from this analysis as dominant strategy equilibria are invariant to risk attitudes. Results are summarized below in Table 6.

Model	**FPSB**	**TPSB**	**All-Pay**	**Literature Range**
CRRA (ρ)	0.54	1.79	0.52	0.3–0.8
CARA (λ)	0.013	1.01	0.01	0.008–0.02

Table 6: **Estimated risk aversion parameters.** Literature range is based on Cox et al. (1988) for FPSB and Dechenaux et al. (2015) for All-Pay auctions. For TPSB, we compare to the limited experimental evidence from Kagel and Levin (1993).

The most robust empirical finding from the current literature is for the FPSB auction; it's a well-known empirical result that bidders in the FPSB bid above the risk-neutral Nash equilibrium level due to risk aversion, thereby failing revenue equivalence (Cox et al., 1988; Kagel and Levin, 1993; Kagel and Roth, 2020). The implied CARA and CRRA parameters align closely with measured risk attitudes in human subjects. For TPSB auctions, estimated risk-aversion parameters are notably larger than both the FPSB parameters and typical literature benchmarks, but still consistent with a coherent risk-averse framework. One plausible explanation for the higher observed risk-aversion in TPSB may be the compounding influence of loss aversion, which biases bids in the same downward direction as risk aversion. The all-pay auction results present a more complicated picture. The estimated parameters for all-pay auctions suggest reasonable risk attitudes. However, these estimates cannot fully explain observed deviations: low-valued bidders, for instance, systematically fail to reduce their bids below equilibrium predictions, contrary to standard risk-aversion reasoning. This inconsistency may indicate a deeper misunderstanding among LLM agents about the strategic complexity inherent to the all-pay mechanism. It is worth noting that human experiments on all-pay auctions also yield highly variable results, reflecting the cognitive challenge this format presents to both human and LLM bidders alike (Dechenaux et al., 2015; Gneezy and Smorodinsky, 2006).

While risk-aversion offers explanatory power for mistakes relative to the risk-neutral benchmark in the

FPSB, TPSB and all-pay auctions, it cannot help explain deviations from truthful bidding in the SPSB. Nevertheless, as with human bidders, LLM bidding in the SPSB auction exhibits numerous departures from equilibrium play. Table 3 summarizes the frequency of underbidding, value bidding, and overbidding by LLM bidders, comparing these results directly with human data from Kagel and Levin (1993). In our experiments, 22% of bids exactly match bidders' values, closely aligning with the 27% reported by Kagel and Levin (1993).[5] Despite similar overall rates of mistakes, the nature of these deviations differs strikingly. In a SPSB auction with 5 players, Kagel and Levin (1993) find the distribution of bids to be right-skewed, with most of the bidding mass (67%) on over-bidding one's value. In contrast, we find the distribution of bids to be primarily left-skewed, with most of the bidding mass (66%) on under-bidding one's value. The contrast highlights that despite both under-bidding and over-bidding being violations of strategy-proofness, they represent two distinct types of behavioral errors empirically – an insight that has not been well rationalized by the auctions literature yet. The prevalence of underbidding among LLM bidders further suggests loss-aversion to a degree not observed in human bidders. And, taken as a whole, we see the mistakes under SPSB as additional motivation for exploring the effect of obviously strategy-proof designs.

In this Section, we've highlighted several features of individual and aggregate LLM bidding behavior. Perhaps the feature of LLM data that makes it the hardest to interpret individual bids is the increased variance. We suggest this increased variance in LLM bidding is due to increased random play by LLMs, due both to a desire to gather exploratory data as well as to disrupt opponents' predictions. The ability to peer into LLMs' chain-of-thought reasoning makes this particularly transparent. For instance, LLM bidders frequently articulate plans like the following from Bidder betty in a round of the FPSB auction:

> "I'll adopt a balanced strategy, bidding 65% of my value. **I'll also introduce random bids occasionally to disrupt predictability**. Monitoring competitor's bids remains essential to adjust my strategy accordingly."

This ability to randomize is just one way in which LLMs may differ from human bidders, but there are many others – say, the ability to reliably calculate 65% of one's value or to monitor and perform inference over the full history of competitor bids. The variance in LLM play suggests to us that looking at LLM data in aggregate, rather than scrutinizing individual data points, may be more productive. Just as individual human bidders encode their own biases and tendencies into their actions, no individual LLM bidder's behavior is particularly instructive – however, in aggregate, auction data with LLM data encodes behavioral patterns that may align closely with many of the empirical regularities from auction theory. It is our view that documenting these points of alignment or divergence will make the future use of LLM data for auction theory or mechanism design more productive.

Finally, one may ask how prompting may change the general direction of our results. This is a first-order concern for the eventual use of synthetic data in mechanism design and experimentation. To address this, we pursue two additional lines of inquiry:

(1) We test prompt robustness. We implement our FPSB and SPSB environments with prompts in a variety of prompt styles, in different languages and with different currencies to study how bids change under prompt variation. These results are reported in Appendix G. We find that naive changes in prompt style, language or currency do not qualitatively change any of the results reported in this Section.

(2) Beyond naive prompt changes, we test prompts which influence the LLM agents' understanding of the underlying economics. To do this, we design prompts to teach LLMs particular mental models or avoid

[5] Note this comparison involves slight measurement differences – Kagel and Levin (1993) allow a small margin (within $\$0.05$) for classifying bids as "truthful," whereas we require exact equality.

particular failure modes; e.g., a prompt to consider Nash deviations, to conceptualize the auction as an extensive-form game, or a prompt which encourages agents to bid in a risk-neutral manner. These results are reported in Section 5. We find that *these* interventions can dramatically change our results.

Our results suggest to us that prompting *can* induce dramatically different behavior in LLMs by making the underlying logic of mechanisms more apparent, but is unlikely to generate different results with naive changes in prompts. It remains an open engineering question as to 'how much' prompting is required to match distributions of existing empirical evidence more closely.

3.2 Obviously Strategy-proofness

For the rest of the paper, we restrict our attention to the second-price auction mechanism. Having considered LLM play in a variety of sealed-bid auctions, we now vary the representation of the second-price auction to make the auction 'easier' to play. Namely, we turn our attention to clock auctions. Clock auctions are dynamic auction formats where the price increases incrementally (or decreases incrementally in reverse auctions), and bidders decide whether to stay in or drop out at each price point. Li (2017) provides an argument as to why clock auctions may be easier to recognize as strategy-proof for humans, and formalizes this as an equilibrium refinement on strategy-proofness called *obvious strategy-proofness* (OSP).

3.2.1 Setting

There are three bidders in each auction but now, following Li (2017), bidders draw *affiliated private values* (APV) of the form $v_i = c + p_i$. The common component c is drawn once, $c \sim U[0, 20]$, and the private component p_i is drawn n times independently, $p_i \sim U[0, 20]$. Winners of the auction receive their own value of the prize v_i when they win, and the 'common' and 'private' components serve to make values correlated. We switch to considering APV settings in this Section for two reasons. First, in spirit of our larger goal to provide many benchmarks of LLM play in various environments. Second, to enable comparison with the existing empirical literature; the two largest experiments of which have been conducted in APV settings.

The *ascending clock auction* (AC) has a global price clock that increases in some fixed increment and bidders who choose when to drop out. The last bidder remaining wins the prize at the price where the second to last bidder dropped out. If multiple bidders are the last to drop out, we resolve ties by randomly allocating the prize among them and setting the winning price to that mutual drop out price. We also study the *blind ascending clock auction* (AC-B), a variant where bidders are not told when other bidders drop out.

3.2.2 Prompts

The sealed-bid prompt follows the style of the last Section's prompts closely. The only difference is in the first two paragraphs, where agents are informed of the APV environment. For the SPSB, we describe the auction format as follows:

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their value for
   the prize. Your value for the prize will be calculated as
   follows:
```

```
1. First we will randomly draw a common value between {{
   common_low}} and {{common_high}}, with all values equally
   likely.
2. Then, for each bidder, a private taste adjustment will be
   drawn between 0 and {{private}}, with all values equally
   likely.
Your value for the prize is equal to the common value plus your
    private taste adjustment. You will not learn the common
   value or your private taste adjustment separately. This
   means that each person in your group may have a different
   value for the prize. However, if you have a high value, it
   is more likely that other people in your group have a high
   value.
After learning your value, you will submit a bid privately at
   the same time as the other bidders. Bids must be between $0
   and ${{common_high + private}} in ${{increment}} increments.
The highest bidder wins the prize and pays the second-highest
   bid. If you win, your earnings will increase by your value
   for the prize, and decrease by the second-highest bid. If
   you don't win, your earnings will remain unchanged.
After each auction, we will display all bids and the winner's
   profits. Ties for the highest bid will be resolved randomly.
```

The implementation of the clock auction simulates a global clock price, with prices beginning at 0 and increasing in $1 increments. The clock auction format descriptions closely mirror the SPSB format above, but with the added description of the clock mechanism in the fourth and fifth paragraphs. In the AC auction, we describe the auction format as follows:

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their value for
   the prize. Your value for the prize will be calculated as
   follows:
1. First we will randomly draw a common value between {{
   common_low}} and {{common_high}}, with all values equally
   likely.
2. Then, for each bidder, a private taste adjustment will be
   drawn between 0 and {{private}}, with all values equally
   likely.
Your value for the prize is equal to the common value plus your
    private taste adjustment. You will not learn the common
   value or your private taste adjustment separately. This
   means that each person in your group may have a different
   value for the prize. However, if 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 at 0
   and counts upwards in {{increment}} USD increments, up to a
   maximum of {{common_high + private}}. At any point, you can
   choose to leave the auction, and anytime a bidder leaves, we
    will broadcast that information to all the remaining
   bidders.
When there is only one bidder left in the auction, that bidder
   will win the prize at the current price. If you win, your
   earnings will increase by your value for the prize, and
   decrease by the current price. If you don't win, your
   earnings will remain unchanged.
After each auction, we will display all bids and the winner's
   profits. Ties for the highest bid will be resolved randomly.
```

The AC-B auction description closely resembles the AC format, but with the explicit caveat that we will not notify bidders when competitors drop out. We describe the auction format as follows:

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their value for
   the prize. Your value for the prize will be calculated as
   follows:
1. First we will randomly draw a common value between {{
   common_low}} and {{common_high}}, with all values equally
   likely.
2. Then, for each bidder, a private taste adjustment will be
   drawn between 0 and {{private}}, with all values equally
   likely.
Your value for the prize is equal to the common value plus your
    private taste adjustment. You will not learn the common
   value or your private taste adjustment separately. This
   means that each person in your group may have a different
   value for the prize. However, if 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 at 0
   and counts upwards in {{increment}} USD increments, up to a
   maximum of {{common_high + private}}. At any point, you can
   choose to leave the auction, but we will not tell any bidder
    when someone leaves.
When there is only one bidder left in the auction, that bidder
   will win the prize at the current price. If you win, your
```

```
    earnings will increase by your value for the prize, and
    decrease by the current price. If you don't win, your
    earnings will remain unchanged.
After each auction, we will display all bids and the winner's
    profits. Ties for the highest bid will be resolved randomly.
```

3.2.3 Theoretical benchmarks

What effect an affiliated values environment has on strategic behavior may not be ex-ante obvious. On one hand, the correlation in values may strengthen incentives for extremal-valued bidders, so that extremal-valued bidders are actually *more* likely to bid their value even in the sealed-bid format (echoing arguments from Li (2017)). On the other hand, APV may be cognitively more complex for bidders relative to IPV environments, thereby inducing errors in play. In either case, the three auction formats in this section are strategically equivalent, and so the affiliation is, in the rational sense, a red herring: for all three auctions it is still a dominant strategy to bid (or drop out at) one's value. The SPSB, AC-B and AC auctions are strategically equivalent and share an equilibrium in dominant strategies:

$$\beta^{**}(v) = v.$$

An agent's bid, in the ascending clock format, is equivalent to the price at which they drop out of the auction at. We will use the terms 'bid' and 'drop-out price' interchangeably in the context of clock auctions. The equilibrium refinement of OSP, introduced by Li (2017), provides a formal framework for analyzing the cognitive simplicity of mechanisms. Consider a strategy S_i for player i. We say that S_i *obviously dominates* another strategy S_i' if, for any deviation from S_i to S_i', the best possible outcome from S_i' is no better than the worst possible outcome from S_i.

Formally, for all histories h where both S_i and S_i' are consistent with h:

$$\min_{z \in Z(S_i,h)} u_i(z) \geq \max_{z' \in Z(S_i',h)} u_i(z')$$

where $Z(S_i, h)$ is the set of terminal histories that can result from playing S_i after history h and $u_i(z)$ is player i's utility for terminal history z. In a clock auction, the history refers to the sequence of events that track bidder activity throughout the auction. Specifically, it records whether and when bidders drop out as the price increases, providing insight into bidding behavior and competition dynamics. A mechanism is said to be obviously strategy-proof if it has an equilibrium in dominant strategies. Both ascending clock auctions studied in this paper are obviously strategy-proof, though the AC-B auction provides bidders with less information than the AC auction by suppressing when bidders drop out.

3.2.4 Empirical benchmarks and results

For each of the clock APV auction formats, we perform 30 experiments, while for the sealed-bid format we perform 10 experiments with 15 rounds. For all experiments, we use $n = 3$ LLM bidders competing over a prize. A single round for the clock format is the run of a global clock from the initial price of \$0 to the final clearing price. The complete simulation of a single clock round is provided in B.2.

Our main results are plotted in Figure 3. We find that the clock formats induce play significantly closer to one's value than the sealed-bid format. Table 3.2.4 summarizes this concisely: the clock formats, aggregated, have an average difference between bid and value of 1.83 (with 73.9% of agents dropping out at their value, excluding winners), while the SPSB bidding has an average difference between bid and value of 6.26 (with

18.7% of all bids equal to an agent's value). This corresponds closely with Li (2017)'s experimental results: aggregated across all rounds, he reports a mean difference between bid and value of 2.48 for the ascending clock auction, and a mean difference of 6.52 for the sealed-bid auction. In general, the clock format improves play for LLM bidders in much the same way as with human bidders.

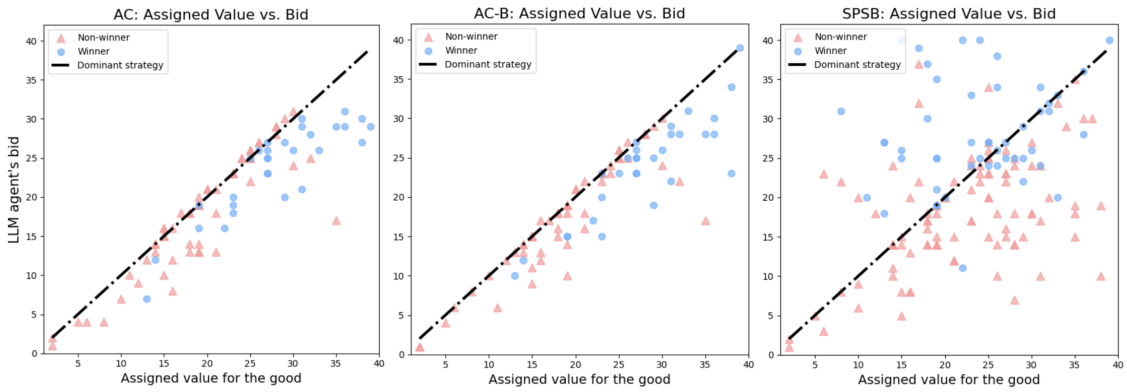


Figure 3: **Comparison of three strategically equivalent auctions (the AC, AC-B and SPSB auctions) in an APV setting.** The AC and AC-B auctions are obviously strategy-proof, whereas SPSB auctions are only strategy-proof. The black dot-dash line represents the dominant strategy where a bidder's bid equals their value. In the clock auctions, red triangles indicate non-winners' drop out prices and blue circles indicate the final price that winners win at. In the SPSB auction, red triangles indicate non-winners' bids and the winning bid in clock auctions. Note: the blue dots are not expected to equal an agents value in the clock auctions, as we do not observe a winner's drop-out price in the clock auction.

Metric	**AC**	**AC-B**	**SPSB**
Truthful Bidding (%)	70.0%	76.7%	18.7%
Average Difference	2.03 (0.39)	1.63 (0.40)	6.26 (0.62)
Sample size	90	90	150

Table 7: **Share of truthful bidding in the AC, AC-B and SPSB auctions in an APV setting**. All bid data comes from non-winners in the clock auctions. We report mean effects, with standard errors reported in parentheses.

Breitmoser and Schweighofer-Kodritsch (2022) provide additional evidence that clock auctions improve truthful play, and also provide evidence that play under AC is slightly better than under AC-B. While we replicate the former result, we fail to replicate the latter result. Play under the AC auction was not significantly different compared to the AC-B auction ($t = 0.71, p = 0.48$) in our experiments. We report all the two-sample t-tests of the mean absolute difference between bids and values in Table 8 – both AC and AC-B mechanisms show significantly smaller bid-value deviations than the SPSB auction ($p < 0.001$), but not from one another.

Comparison	t-statistic	p-value
AC v.s. AC-B	0.71	0.48
AC v.s. SPSB	-6.28	3.21e-09
AC-B v.s. SPSB	-5.78	4.09e-08

Table 8: **Two-sample t-tests comparing the mean absolute difference between bids and values in the AC, AC-B and SPSB auctions in an APV setting**. Results demonstrate that both clock auctions have lower mean absolute difference between bid and value vs the sealed-bid auction. However, the clock auctions themselves are not significantly different. Bid data for clock auctions is excluding winners.

Finally, we observe that the clock format improves on a particular failure mode of the SPSB—namely, bidding *above* one's value. The result is that the share of total mistakes in clock auctions overwhelmingly comes from dropping out too early rather than too late. Although OSP mechanisms cannot rationalize a difference between over- and under-bidding (both represent deviations from an obviously dominant strategy), this evidence further suggests that over- and under-bidding represent categorically *distinct* failure modes. Further exploration into rationalizations of the difference between these two failure modes remains an interesting and open question for future research.

3.3 Winner's curse

The last set of simulations we run are in the common value setting. We demonstrate evidence for the winner's curse with $n = 2, ..., 6$ bidders following the classic paper from Kagel and Levin (1986).

3.3.1 Setting

In the *common value* setting, there are n bidders varying from $n = 2 \dots 6$. Bidders draw values of the form $v_i = c + p_i$. The common component is sampled once, $c \sim U[20, 79]$, and the private shock component p_i is sampled n times independently, $p_i \sim U[-20, 20]$, with p the vector of all private shocks. In the common value setting, every bidder's true value for the good is c but each bidder has a noisy estimate of this value. Hence, agents bid based on values $v_i = c + p_i \in [0, 99]$ with a trapezoidal distribution, but only obtain c when they win, $u_i(\beta(v)) = \mathbb{1}_{i \text{ won the auction}} \cdot (c - \beta^{(2)}(v))$. The auction is run as an SPSB auction to minimize the effect of strategic shading (as in the BNE of the FPSB auction) with uniform values as much as possible.

3.3.2 Prompts

The prompts here follow the style of the previous prompts closely. The description of the common value environment follows the language of the affiliated environment closely, while description of the allocation and transfer rule remain unchanged.

For the SPSB auction with common values, we describe the auction format as follows:

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their perceived
   value for the prize - a noisy measurement of the true value
   of the prize. Your perceived value for the prize will be
   calculated as follows:
```

```
1. For each round, a common value will be drawn between {{
   common_low}} and {{common_high}}, with all values equally
   likely to be drawn.
2. For each person, a private noisy adjustment will be drawn
   between -{{private}} and {{private}}, with all values
   equally likely to be drawn.
We will tell you your perceived value, the sum of the common
   value and the private noise adjustment. However, everyone's
   true value for the prize is equal to the shared common value
   .
After learning your perceived value, you will submit a bid
   privately at the same time as the other bidders. Bids must
   be between $0 and ${{common_high + private}} in ${{increment
   }} increments.
The highest bidder wins the prize and pays the second-highest
   bid. If you win, your earnings will increase by the true
   value for the prize, and decrease by the second-highest bid.
    If you don't win, your earnings will remain unchanged.
After each auction, we will display all bids and the winner's
   profits. Ties for the highest bid will be resolved randomly.
```

3.3.3 Theoretical benchmarks

Naively, under the SPSB design, a bidder wants to bid their valuation for the good given their information, $\beta_i(v) = \mathbb{E}[c|v_i] = v_i - \mathbb{E}[p_i] = v_i$ if $v_i \in [40, 59]$, else shading linearly to reflect that the entire support of the shock is not possible (i.e., if an agent draws $v_i = 99$, they can immediately infer $c_i = 79$). [6]

However, this naive bid overbids relative to the BNE of the common value auction game, as it neglects to consider adverse selection: the winner i in a symmetric, increasing BNE is the bidder who obtained the highest private shock signal, $p_i = p^{(1)}$. This overbidding is familiar as the "winner's curse."

Hence, the theoretical optimum is for bidders to condition on both events and bid $\beta_i(v) = \mathbb{E}[c \,|\, v_i \wedge p_i = p^{(1)}]$; i.e., conditioning on their value and on the event that their received private shock was the highest (the second event is equivalent to the event that the bidder won with strictly monotone increasing $\beta(\cdot)$).

3.3.4 Empirical benchmarks

The primary experimental literature makes two positive predictions on the common value auction (Kagel and Levin, 1986). First, even experienced bidders fail to condition on the event where their signal is the highest (called 'item valuation considerations' in their paper), thereby still falling victim to the winner's curse; second, the winner's curse barely shows up in small auctions (3-4 bidders) but has an effect in big auctions (6-7 bidders).

[6] More precisely, when $v_i \in [40, 59]$ the full support $[-20, 20]$ applies so $E[p_i \mid v_i] = 0$ and $E[c \mid v_i] = v_i$. For $v_i < 40$, the constraint $c \geq 20$ forces p_i to lie in $[-20,\ v_i - 20]$ with mean $\frac{v_i - 40}{2}$, yielding $E[c \mid v_i] = \frac{v_i + 40}{2}$; similarly, for $v_i > 59$ the constraint $c \leq 79$ forces p_i into $[v_i - 79,\ 20]$ with mean $\frac{v_i - 59}{2}$, so $E[c \mid v_i] = \frac{v_i + 59}{2}$.

3.3.5 Simulation evidence

We find evidence corroborating each of these predictions. In auctions of all sizes, bidders do not shade up or shade down their bids (i.e., respecting that $\mathbb{E}[p] = 0$)) but fail to realize that if they win, 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 4.

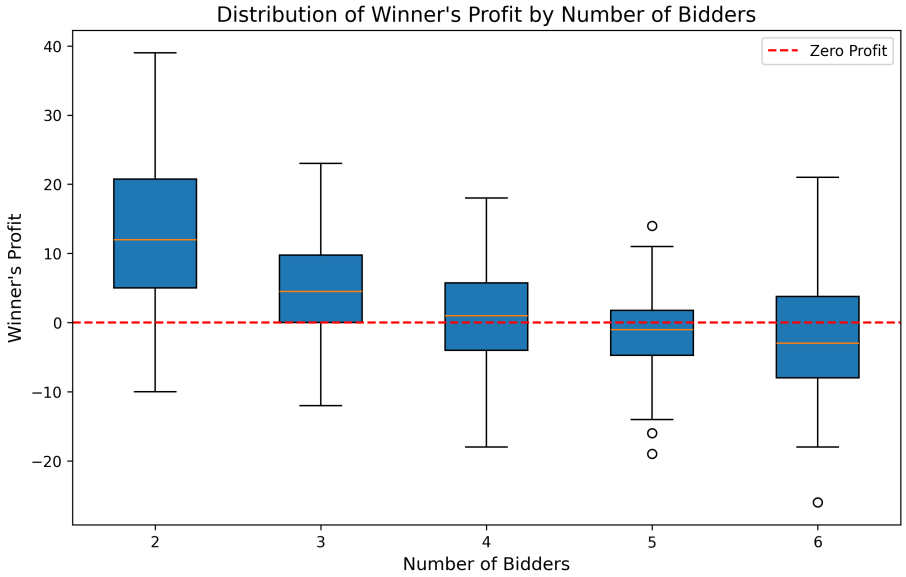


Figure 4: **Distribution of the winner's profit across SPSB auctions with 2 to 6 bidders in a CV setting.** Each box shows the interquartile range of profits, with the median indicated by the central yellow line. The horizontal red dashed line represents zero profit. As the number of bidders increases, the median winner's total profit decreases and more frequently turns negative, illustrating the intensifying effect of the winner's curse in larger auctions. Circles are extremal values outside of the 95% CI.

While play in a non-common value environment for the SPSB is invariant to the number of bidders, with common values, the bite of adverse selection hurts the winner more with more bidders. Our evidence suggests that LLMs do not learn to appreciate this bite, even as the cost of adverse selection increases. This mirrors results with humans, where even experienced bidders struggle to update bidding to avoid the winner's curse. One difference between our results and human results is that even with many rounds LLMs seem not to figure out how to respond to the winner's curse. A defense of LLMs here may be that this is just very hard: learning is non-trivial with human agents as well (in Kagel and Levin (1986) it takes players 15-20 periods to learn), and in some studies humans never learn (i.e., in the bilateral bargaining games of **?**). This suggests more sophisticated learning techniques may be required to fully mitigate the winner's curse with LLM agents.

4 Ebay Auction Design

Having tested LLMs against classic results in auction theory, we build an auction environment to test components of auction design in a setting inspired by eBay's online marketplace.

eBay's auction platform, studied by Einav et al. (2011) and others, features various design elements—for example, a hidden reserve price (as opposed to a reserve price) and proxy bidding—that are *in theory* less critical to auction outcomes yet *in practice* appear to have substantial effects on bidder behavior. These differences highlight the gap between textbook theory and real-world bidding, where seemingly small design details can alter bidder behavior and thus price discovery and revenue.

In this section, we simulate a stylized version of the eBay auction design to consider two design parameters more closely: the hidden reserve price and the implementation of a 'closing rule'. The core mechanics of an eBay auction are simple: bidders on eBay use a proxy bidding system, which means that each bidder submits their "maximum bid" to the platform, and the platform automatically bids on their behalf, only as much as necessary to keep the winning bidder in the lead (i.e., the second-highest bid + \$1, with incremental bidding managed automatically). Bidders can increase their maximum bid at any point in time, until some closing time of the auction, giving this the flavor of an incremental second-price auction.

The use of reserve prices in auctions is well understood in auction theory. Notably, Myerson (1981) showed that setting an optimal reserve price can maximize expected revenue in certain private value settings (for "regular" distributions). However, eBay also allows sellers (for a fee) to set a *hidden* reserve price: a threshold that bidders do not know until it is met by the highest maximum bid. While there is some work suggesting a hidden reserve can increase expected revenue in common value settings (e.g., Vincent (1995)), in a private values environment with rational bidders using proxy bidding, the hidden reserve should not change behavior relative to a known reserve or even no reserve: bidders should still submit their maximum value upfront. While it is impossible to say whether bidders are bidding their true values, a large empirical study of hidden reserves on eBay suggests that bidders do not change bidding behavior much in the face of hidden vs. non-hidden reserves (Einav et al., 2011).

A more dramatic departure from simple, truthful bidding is 'bid sniping,' or the practice of placing bids in the final moments of an auction to avoid triggering competitive responses. Sniping is prevalent on eBay, with dedicated e-snipe services automating last-second bid submissions.[7] This phenomenon has been studied for decades, with explanations ranging from best responses to irrational types to behavioral inertia (Roth and Ockenfels, 2002). Such wide prevalence of sniping may be especially troubling to economists, as it inhibits a core function of auctions: their ability to induce timely, efficient price discovery. One avenue to curb sniping may be a 'closing rule' (also called a 'soft-close' or 'auction extension' rule)—namely, when a new maximum bid comes in at the last minute, extend the auction close to allow for more bidding. This design detail has been credited with the substantial difference in bidding dynamics between earlier auctions on Amazon.com, which had this auction-extension rule, and auctions on eBay, which do not. Specifically, Roth and Ockenfels (2002) and Ockenfels and Roth (2006) show that much less sniping occured on Amazon.com, with bids on Amazon.com occuring throughout the auction compared to the empirical evidence of sniping on eBay.

4.1 Setting

In each eBay simulation, there are three bidders. Without the modified closing rule, the auction runs for 10 bidding periods. Bidders draw private values uniformly and independently from $[0, 99]$. In regard to the auction environment, we consider four treatments, varying the hidden reserve and the closing rule:

[7] E.g., `https://www.esnipe.com/`.

T1 : Standard eBay auction with proxy bidding.

T2 : eBay auction with a modified (auction extension) closing rule.

T3 : Standard eBay auction with a hidden reserve price.

T4 : eBay with a modified (auction extension) closing rule and a hidden reserve price.

For each hidden reserve price design, we run experiments with hidden reserves at $r = 40, 50$, and 60.

We model continuous time in the synthetic, modified eBay auction environment as follows: each period, bidders bid (i.e., decide whether to increase their maximum bid or hold) in a common knowledge, randomly-permuted order (e.g., today Charles, then Alice, then Betty will bid; tomorrow, Alice, then Betty, then Charles). Moreover, each bidder can decide whether to increase their maximum bid with knowledge of the updated state of the auction based on the actions of bidders that precede them that period. In the final period, bidding is simultaneous and no bidder will know whether they are the last bidder (simulating that, in a continuous time auction, you cannot know if another bidder will get a bid in between you and the auction's close). Under the modified closing rule, if a higher bid is placed in the final period such that there is a new winner, the auction is extended by an additional period. This process continues until no further bids are made.

With this stylized eBay framework we abstract away from additional eBay design features such as shipping fees and Buy-It-Now. There was no prompting of LLMs to induce sophisticated play as in Section 5. The complete simulation process is explained in Appendix B.3 and the prompts are listed in Appendix G.

4.2 Simulation Process of eBay Auction

For eBay auctions, we discretize the continuous bidding period into 10 periods. Each period, LLM agents decide whether to increase their bid or hold their current bid. The bidding process follows a structured format, where bidders act in a predefined sequence each period (e.g., Charles, then Alice, then Betty). For each round, bidders are reminded of the auction rules, similar to those outlined in Section 2.2. Agents are informed of past price changes through a transcript, such as "On day 1, the price changed to 1. On day 2, the price changed to 3." The detailed prompt and an example of one clock are listed in B.3.

4.3 Simulation evidence

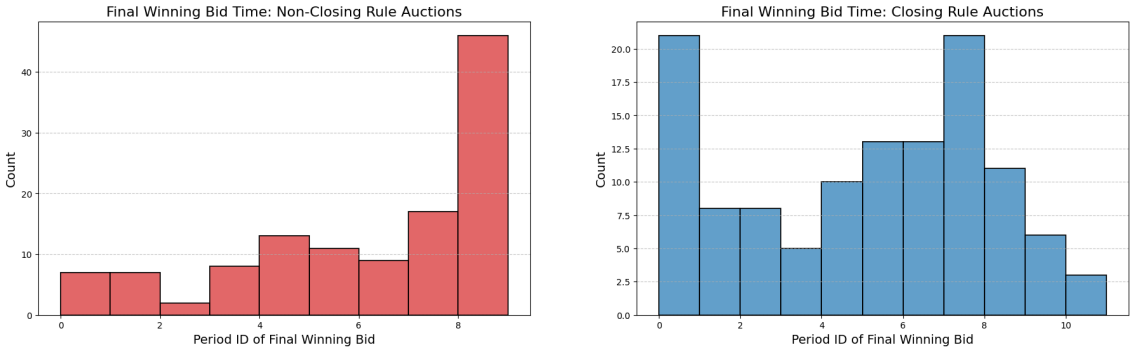
Hidden reserve price results.

Figure 13 shows that the auction revenue remains basically unchanged regardless of the experiment type. In Table 10, we report a table of two-sided t-tests making this explicit.

Closing rule results. Second, we demonstrate that modifying eBay's closing rule to allow for an extension of the auction when there is final-period bidding has a major effect on bidder behavior. For this, we say the 'final winning bid time' is the period at which the winner is settled; i.e., the period in which the maximum bid was placed by the winning bidder.[8] This serves as a measure of when the winning bidder is determined. Figure 5 reports the 'final winning bid time' as a histogram across auction instances with the modified closing rule and without the modified closing rule. Figure 6 reports a CDF across the four treatment types (T1 - T4). Without the modified closing rule, we see that LLMs predominantly bid in the period before close; however, with the modified closing rule, this behavior disappears. In Figure 6, we see the same behavior. With

[8]Note this means that if a winning bidder increases their maximum bid after already being the maximum bidder, winning bid time does not update (as no additional discovery was conducted through the auction).

the modified closing rule, the proportion of last-period bids drops substantially; adding a hidden reserve also reduces this behavior. This latter effect is somewhat surprising, but one possible explanation is that a hidden reserve increases bidder uncertainty, discouraging last-moment bidding and shifting bids earlier in the auction, as suggested by Bajari and Hortacsu (2003). In effect, the modified closing rule shifts the winning bid time left on the histogram, so that the auction's winner is chosen earlier.



(a) Non-Closing Rule

(b) Closing Rule

Figure 5: **Histogram of final winning bid in IPV ebay-style auctions with a standard eBay closing rule vs. modified eBay closing rule.** The modified closing rule introduces the auction-extension logic, moving the timing of the final winning bid substantially earlier in the auction.

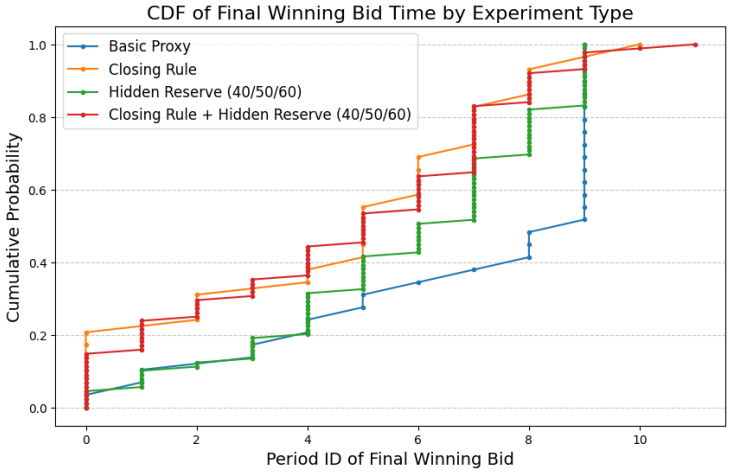


Figure 6: **CDF of final winning bid in IPV eBay-Style auctions.** The cumulative distribution function of the final winning bid time for each auction format. Formats with a modified closing rule end earlier (with and without a hidden reserve), especially compared to the base design without a hidden reserve.

5 Intervention design in Sealed-Price Auction

It matters how agents perceive mechanisms. One seminal example from the market design literature is that of the design of the *National Residency Match Program* (NRMP) for doctors. Since the early days of the NRMP, students have mistakenly strategically reported their preferences to the match program despite it being a strategy-proof matching mechanism, thereby hurting the overall efficiency of the system (Rees-Jones, 2018; Roth and Peranson, 1999). Furthermore, improving our understanding of agents' perception of mechanisms may also improve our ability to design simple mechanisms Li (2024).

This section introduces experimental interventions that seek to influence how agents understand a mechanism. We begin with three conceptual mechanisms that offer different frameworks for determining the equilibrium logic of an auction. We then provide two interventions that probe the effect of giving advice. We also provide one behavioral intervention which seeks to affect bidders' risk preferences.

In the literature, Gonczarowski et al. (2023) present a method to highlight and emphasize strategy-proofness with short, one-sentence descriptions in an extensive-form menu[9]. However, their menu-description treatment of a second-price sealed-bid (SPSB) auction did not perform well in the field when they tested it empirically. They argue this may be because the classic descriptions of some mechanisms may be as transparent as the extensive-form menu description. Breitmoser and Schweighofer-Kodritsch (2022) show that re-framing a static sealed-bid auction as an ascending clock auction can improve the rate of truth-telling behavior, and suggest merely *considering* a clock price may improve play. Finally, the logic of Nash deviations is a classic one to economics and often used to teach contingent thinking.

Here, we present six interventions on the SPSB auction. Of the six interventions, the first three suggest to agents particular conceptual frameworks, the next two provide (good and bad) advice, and the last is a behavioral intervention.

- **Menu-Description (menu_intervention)**: A treatment providing the extensive form of the auction as a menu to the agent.
- **Clock-Description (proxy_intervention)**: A treatment describing the ascending-clock process. Though the proxy clock is described to the agent before the game begins, the auction is still operated as an SPSB auction.
- **Nash-Deviation (nash_intervention)**: A treatment providing an example of a Nash-style thought process; e.g., What happens if you bid up? What happens if you bid down?
- **Direct Revelation (dominant_strat_intervention)**: A treatment directly revealing the dominant strategy of the auction as that of bidding one's value.
- **Wrong Direct Revelation (wrong_strat_intervention)**: An ablation study incorrectly revealing the dominant strategy of the auction as bidding *half* one's value.
- **Risk-neutrality prompting (risk_intervention)**: A behavioral treatment defining risk-neutrality in terms of gambles and then prompting the LLM to be more risk-neutral.

The full intervention prompts are provided in Appendix H and are appended to the end of the rule explanation prompt when eliciting LLM bids.

These interventions aim to reveal which framing or description of the auction design encourages more straightforward bidding strategies among LLMs. Because LLM-based experimentation is relatively low-cost

[9] An extensive-form menu refers to a structured presentation of options or strategies. It outlines the sequential choices available to players, detailing how each decision can lead to subsequent options and outcomes.

and can be repeated rapidly, this provides a useful testbed for exploring new interventions before conducting costly human-subject experiments.

5.1 Simulation evidence

The results of the six interventions are reported in Figure 7. To characterize the error from the theoretical prediction at `bid = value`, we also report the R^2 against $\beta^{**}(v) = v$ in Table 9. However, as was clear in the discussion of the empirical literature in Section 3.1, overbidding and underbidding constitute different failure modes—in particular, many human bidders suffer the failure mode of overbidding in the SPSB auction.

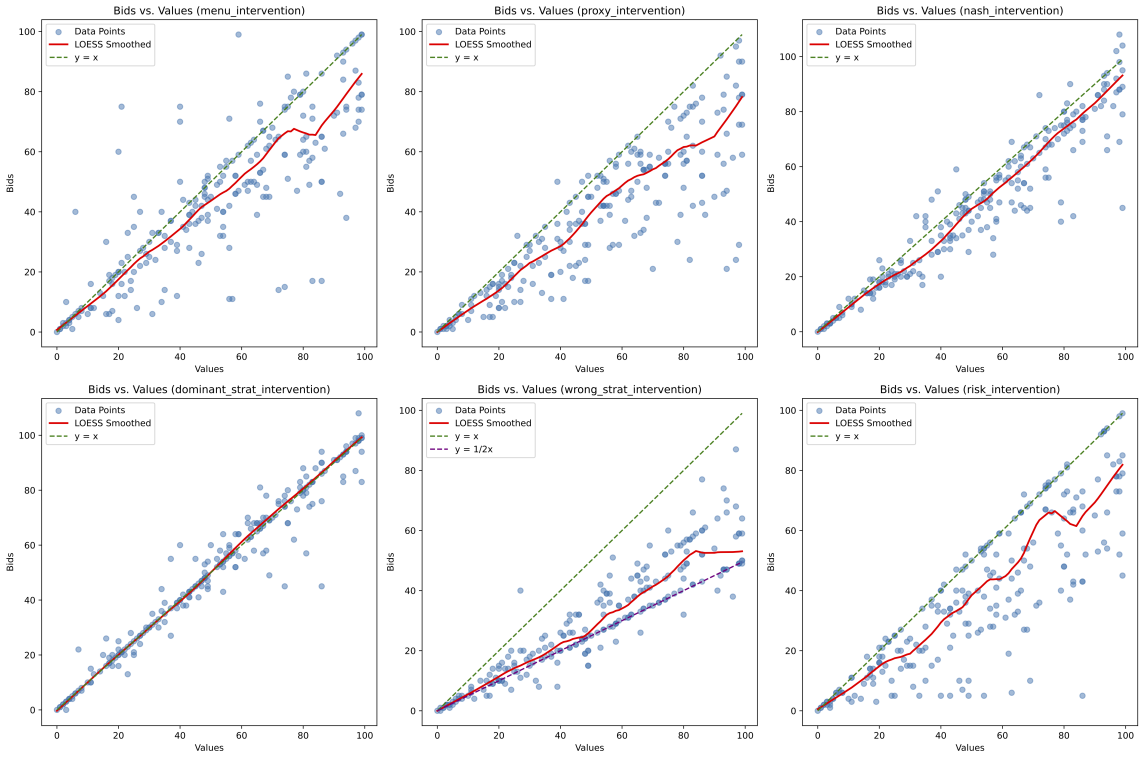


Figure 7: **Results for SPSB with the six interventions, IPV environment.** The green dotted line represents the scenario of bid=value, blue dots are bids from agents, and the red line is the LOESS smoothed curve of bids. The top 3 interventions show the effect of suggesting particular conceptual frameworks to agents. The menu intervention has a null effect. The proxy clock intervention completely eliminates overbidding in the SPSB, but still has underbidding. The Nash intervention dramatically improves along both overbidding and underbidding. The first 2 interventions of the bottom row show that agents follow advice, but not terminally if the advice is bad. The final, risk-neutrality intervention does not succeed in making agents play better.

As such, we also decompose the R^2 into the sum of squares above and below $\beta^{**}(v) = v$ to report each intervention's performance in terms of the two components of overbidding and underbidding. We report

	Intervention Type	R^2	SS Above	SS Below	Prop Above	Prop Below
	No Intervention	0.4845	7,213	78,652	0.0840	0.9160
Conceptual Framework	Menu	0.5337	11,104	57,685	0.1614	0.8386
	Proxy	0.2966	177	85,488	0.0021	0.9979
	Nash	0.8365	1,160	24,591	0.0450	0.9550
Advice	Dominant Strategy	0.9533	2,665	5,521	0.3256	0.6744
	Wrong Strategy	-0.8258	169	141,137	0.0012	0.9988
Behavioral	Risk	0.4004	59	89,920	0.0007	0.9993

Table 9: **Intervention Type Analysis by Category, SPSB and IPV environment.** The R^2 of simulation evidence on the dominant strategy prediction $\beta^{**}(v) = v$. In the latter four columns we report the sum of squares and the proportion of bids above and below the predicted line, respectively. Of the conceptual frameworks, the Nash deviations improve play the most. The advice interventions show that LLMs follow good advice but are unconvinced with bad advice. The behavioral intervention fails to teach risk-neutrality to current LLM agents.

each of these interventions alongside the results from the original SPSB experiments of Section 3.1 (No Intervention) in Table 9. This provides useful descriptive statistics for the effect of the interventions. The first three interventions—menu, proxy and Nash deviation—are meant to teach the LLMs particular ways of reasoning through the problem.

The menu intervention modestly improves on the basic SPSB by reducing underbidding, but has the weakest effect of these three 'anchoring' interventions. There is little effect on either type of failure mode, and the menu intervention, qualitatively, seems to make overbidding a bigger problem. In contrast, the proxy almost completely eliminates the overbidding failure mode. While model fit to $\beta^{**}(v) = v$ goes slightly down because underbidding is worse, our simulations reveal that the proxy intervention improves LLM play by decreasing the overbidding failure mode.

The Nash deviation intervention has a more positive effect than the proxy intervention. Instructing LLMs to reason through the Nash deviations from their chosen bidding strategy dramatically improves performance, with very similar gains on overbidding as with the proxy intervention and decreasing the underbidding failure mode.

The next two interventions feature advice from a trusted source. Of particular note here is the ablation study that offers models the *wrong* advice—there is a concern in experimental economics in regard to "polluting subject pools" by lying to them. That is not a constraint with LLMs, enabling a wider set of possible interventions.[10]

As expected, revealing the dominant strategy to LLMs induces fairly tight play to the dominant strategy $\beta^{**}(v) = v$. Inspecting planning narratives reveals that the overbidding and underbidding errors correspond to LLMs 'testing' the economists recommendation by slightly deviating. Interestingly, the 'wrong strategy' intervention induces the worst play of all six interventions. However, it does so in an interesting way. Though many LLMs follow the recommendation of bidding $\beta(v) = \frac{v}{2}$, they cease to after early rounds of negative profit. However, they are unable to figure out the correct strategy upon being anchored incorrectly, so remain

[10]This is at least true with the current generation of LLMs which are not permanently learning from prompting.

stuck underbidding.

Finally, we implement a single behavioral intervention, called so because risk preferences are a behavioral feature. The intervention has about the same performance as the proxy intervention: instructing LLMs to be risk-neutral has the effect of removing over-bidding. While this intervention improves efficiency, we are primarily optimistic about this behavioral intervention as a proof-of-concept: we hope that future work can characterize more precisely *how* prompting induces various behavioral postures in agents.

6 Discussion

6.1 Design principles and simulation paradigms

Simulations are proliferating in high-stakes engineering disciplines, where iterative design and validation are performed long before prototypes are built. In contrast, traditional mechanism design research often relies on expensive, time-intensive human subject experiments or confounded observational data. Our synthetic lab leverages the rapidly advancing capabilities of large language models to simulate bidder behavior at a fraction of the cost. By systematically varying auction parameters, rules, and descriptive framings in a controlled setting, we not only replicate key empirical regularities but also provide new data for why subtle design differences can yield divergent outcomes in practice. In doing so, this simulation paradigm complements traditional analytical models by integrating behavioral factors into mechanism design.

LLM agents operate on a text-based reasoning framework that captures nuanced decision-making processes. This also affords the opportunity to query these LLM agents about their internal deliberations, opening up the prospect, with further study and validation, that LLMs may also provide information on the human cognitive constraints that underlie deviations from classical auction theory.

6.2 The advantage of LLM participants

Relative to Traditional Lab Experiments Classical lab experiments in auction settings have yielded valuable insights, but also come with high cost and logistical challenges. In stark contrast, our experiments with LLM agents reduce these costs by roughly three orders of magnitude (e.g., from over $\$15{,}000$ to $\$10$ per experiment (Li, 2017)). This dramatic improvement in cost efficiency enables rapid iteration, allowing researchers to explore a broader parameter space with minimal investment of money and time.

Relative to Observational Data The analysis of observational datasets from platforms like eBay is often challenging because one needs to control for the heterogeneity among sellers and buyers, this complicating causal inference. In our synthetic lab, we can isolate individual auction parameters and rule framings by carefully controlling the environment in which bids are generated including the simulated goods and valuations of bidders. This controlled setting is perfect for causal inference and enables us to directly link subtle changes in auction design to variations in bidder behavior (Manning et al., 2024).

Online Laboratory Framework The rise of online human labs (Horton et al., 2011) in behavioral economics has previously lowered the barrier to performing controlled experiments (Enke et al., 2024; Oprea, 2024), enabling near real-time data collection and hypothesis testing. Our synthetic lab extends this idea by eliminating many of the ethical, logistical, and opportunity-cost issues associated with human experiments. By using LLM agents, we gain the flexibility to incorporate additional modules (such as reasoning queries) into the experimental design without risking harm to human subjects. This approach accelerates research and also offers a platform for cross-validating findings with more traditional online lab experiments.

6.3 Comparison with Traditional Auction Test Suites

Existing tools such as the *Combinatorial Auction Test Suite (CATS)* (Leyton-Brown et al., 2000) focus on providing valuation distributions on which to evaluate the efficiency and computational performance of winner-determination and pricing algorithms. While these suites provide a benchmark set of instances with which to study algorithmic performance and auction design they do not seek to model bidder behavior. In this sense, the use of LLMs as proxy models for human bidders complements these test suites, offering a means to probe the behavioral dimensions bidder behaviors in auction design.

7 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 from labs involving human subjects (e.g., evidence of risk-averse bidding and evidence that clock auctions are easier to play). The results are encouraging and we see the main contribution of this work as putting forward a framework for how to think about LLM experimental agents as a proxy for human agents. In particular, the design space for prompting is large, and we hope that future research will make use of our code to run simulations testing their own prompt variations.

While this paper focuses on auction design, future work may use LLM sandboxes to test other kinds of economic mechanisms; e.g., voting, matching, contract design, etc. As LLM models are increasingly validated as proxies for human behavior, LLM agents can then 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.

We are particularly interested in the use of LLM-based techniques to generate synthetic data that is useful for informing economic design. In particular, some auction formats, such as combinatorial auctions, are complex and difficult to run frequently and at scale in traditional lab experiments. Augmenting these traditional lab experiments with LLM-agent experiments, when correctly validated, may open up wide new avenues to better understand the design tradeoffs in these kinds of complex and often high-stakes environments.

References

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

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

Henry Bae, Aghyad Deeb, Alex Fleury, and Kehang Zhu. 2023. ComplexityNet: Increasing LLM Inference Efficiency by Learning Task Complexity. *arXiv preprint arXiv:2312.11511* (2023).

Max H Bazerman and William F Samuelson. 1983. I won the auction but don't want the prize. *Journal of conflict resolution* 27, 4 (1983), 618–634.

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

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

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

Edward C Capen, Robert V Clapp, and William M Campbell. 1971. Competitive bidding in high-risk situations. *Journal of petroleum technology* 23, 06 (1971), 641–653.

Gary Charness and Dan Levin. 2009. The origin of the winner's curse: a laboratory study. *American Economic Journal: Microeconomics* 1, 1 (2009), 207–236.

Jiangjie Chen, Siyu Yuan, Rong Ye, Bodhisattwa Prasad Majumder, and Kyle Richardson. 2023. 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).

Vicki M Coppinger, Vernon L Smith, and Jon A Titus. 1980. incentives and behavior in english, dutch and sealed-bid auctions. *Economic inquiry* 18, 1 (1980), 1–22.

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

Emmanuel Dechenaux, Dan Kovenock, and Roman M Sheremeta. 2015. A survey of experimental research on contests, all-pay auctions and tournaments. *Experimental Economics* 18 (2015), 609–669.

Liran Einav, Theresa Kuchler, Jonathan D Levin, and Neel Sundaresan. 2011. *Learning from seller experiments in online markets*. Technical Report. National Bureau of Economic Research.

Benjamin Enke, Thomas Graeber, Ryan Oprea, and Jeffrey Yang. 2024. *Behavioral attenuation*. Technical Report. National Bureau of Economic Research.

Sara Fish, Yannai A Gonczarowski, and Ran I Shorrer. 2024. Algorithmic Collusion by Large Language Models. *arXiv preprint arXiv:2404.00806* (2024).

Uri Gneezy and Rann Smorodinsky. 2006. All-pay auctions—an experimental study. *Journal of Economic Behavior & Organization* 61, 2 (2006), 255–275.

Yannai A Gonczarowski, Ori Heffetz, and Clayton Thomas. 2023. *Strategyproofness-exposing mechanism descriptions*. Technical Report. National Bureau of Economic Research.

John Horton, Apostolos Filippas, and Robin Horton. 2024. *EDSL: Expected Parrot Domain Specific Language for AI Powered Social Science*. Whitepaper. Expected Parrot.

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

John J Horton, David G Rand, and Richard J Zeckhauser. 2011. The online laboratory: Conducting experiments in a real labor market. *Experimental economics* 14 (2011), 399–425.

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

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

John H Kagel and Dan Levin. 1993. Independent private value auctions: Bidder behaviour in first-, second- and third-price auctions with varying numbers of bidders. *The Economic Journal* 103, 419 (1993), 868–879.

John H Kagel and Alvin E Roth. 2020. *The handbook of experimental economics, volume 2*. Princeton university press.

Vijay Krishna. 2009. *Auction theory*. Academic press.

Kevin Leyton-Brown, Mark Pearson, and Yoav Shoham. 2000. Towards a universal test suite for combinatorial auction algorithms. In *Proceedings of the 2nd ACM conference on Electronic commerce*. 66–76.

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

Shengwu Li. 2024. Designing Simple Mechanisms. *Journal of Economic Perspectives* 38, 4 (2024), 175–192.

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

Paul Milgrom and Ilya Segal. 2020. Clock auctions and radio spectrum reallocation. *Journal of Political Economy* 128, 1 (2020), 1–31.

Dov Monderer and Moshe Tennenholtz. 2000. K-price auctions. *Games and Economic Behavior* 31, 2 (2000), 220–244.

Roger B Myerson. 1981. Optimal auction design. *Mathematics of operations research* 6, 1 (1981), 58–73.

Axel Ockenfels and Alvin E Roth. 2006. Late and multiple bidding in second price Internet auctions: Theory and evidence concerning different rules for ending an auction. *Games and Economic behavior* 55, 2 (2006), 297–320.

Ryan Oprea. 2024. Decisions under risk are decisions under complexity. *American Economic Review* 114, 12 (2024), 3789–3811.

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

Pratyush Patel, Esha Choukse, Chaojie Zhang, Aashaka Shah, Íñigo Goiri, Saeed Maleki, and Ricardo Bianchini. 2023. Splitwise: Efficient generative llm inference using phase splitting. *Power* 400, 700W (2023), 1–75.

Narun Raman, Taylor Lundy, Samuel Amouyal, Yoav Levine, Kevin Leyton-Brown, and Moshe Tennenholtz. 2024. STEER: Assessing the Economic Rationality of Large Language Models. arXiv:2402.09552 [cs.CL] https://arxiv.org/abs/2402.09552

Alex Rees-Jones. 2018. Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match. *Games and Economic Behavior* 108 (2018), 317–330.

Alvin E Roth and Axel Ockenfels. 2002. Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet. *American economic review* 92, 4 (2002), 1093–1103.

Alvin E Roth and Elliott Peranson. 1999. The redesign of the matching market for American physicians: Some engineering aspects of economic design. *American economic review* 89, 4 (1999), 748–780.

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

Daniel R Vincent. 1995. Bidding off the wall: Why reserve prices may be kept secret. *Journal of Economic theory* 65, 2 (1995), 575–584.

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

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* 35 (2022), 24824–24837.

A Proofs of auction equilibria

A.1 Risk-neutral equilibria

First, we prove the risk-neutral cases.

A.2 Risk-neutral equilibria

In this section, we derive the Bayesian Nash equilibria (BNE) for different auction formats under risk neutrality. We begin with the second-price sealed-bid auction, then derive the equilibria for first-price, third-price, and all-pay auctions.

A.2.1 Second-price sealed-bid auction (SPSB) dominant strategy equilibrium

In a second-price sealed-bid auction with independent private values drawn from a common distribution F, bidding one's true valuation $B(v) = v$ is a weakly dominant strategy.

Proof. This follows directly from Vickrey (1961).

A.2.2 First-price sealed-bid auction (FPSB) risk-neutral BNE

In a first-price sealed-bid auction with n risk-neutral bidders whose valuations are independently drawn from a uniform distribution on $[0, \bar{v}]$, the symmetric Bayesian Nash equilibrium bidding strategy is $B(v) = \frac{n-1}{n}v$.

Proof. Consider a bidder with valuation v who bids b. Suppose all other bidders follow a strictly increasing bidding strategy $B(v)$. The bidder wins if and only if $b > B(v_j)$ for all other bidders j, which occurs with probability $[F(B^{-1}(b))]^{n-1}$ where F is the distribution function of valuations.

The expected payoff for this bidder is:

$$\pi(b, v) = (v - b)[F(B^{-1}(b))]^{n-1} \tag{1}$$

For a uniform distribution on $[0, \bar{v}]$, we have $F(v) = \frac{v}{\bar{v}}$. Taking the first-order condition with respect to b and imposing the equilibrium condition $b = B(v)$, we get:

$$\frac{\partial \pi}{\partial b} = -(F(B^{-1}(b)))^{n-1} + (v - b)(n - 1)(F(B^{-1}(b)))^{n-2} f(B^{-1}(b)) \frac{d}{db} B^{-1}(b) = 0 \tag{2}$$

In a symmetric equilibrium, $B^{-1}(b) = v$, so $\frac{d}{db}B^{-1}(b) = \frac{1}{B'(v)}$. With the uniform distribution, $f(v) = \frac{1}{\bar{v}}$. Substituting these into the FOC and simplifying:

$$-\left(\frac{v}{\bar{v}}\right)^{n-1} + (v - B(v))(n - 1)\left(\frac{v}{\bar{v}}\right)^{n-2} \frac{1}{\bar{v}} \frac{1}{B'(v)} = 0 \tag{3}$$

This simplifies to:

$$-v^{n-1} + (v - B(v))(n - 1)v^{n-2} \frac{1}{B'(v)} = 0 \tag{4}$$

Rearranging:

$$B'(v) = \frac{(n - 1)v^{n-2}(v - B(v))}{v^{n-1}} = \frac{(n - 1)(v - B(v))}{v} \tag{5}$$

This is a first-order differential equation. To solve it, we set $B(v) = vR(v)$ for some function $R(v)$. Then $B'(v) = R(v) + vR'(v)$. Substituting:

$$R(v) + vR'(v) = \frac{(n-1)(v - vR(v))}{v} = (n-1)(1 - R(v)) \tag{6}$$

Simplifying:

$$R(v) + vR'(v) = (n-1) - (n-1)R(v) \tag{7}$$

Rearranging:

$$vR'(v) = (n-1) - nR(v) \tag{8}$$

This is a linear first-order differential equation. The general solution is:

$$R(v) = \frac{n-1}{n} + \frac{C}{v^n} \tag{9}$$

where C is a constant of integration. For the boundary condition, we know that a bidder with valuation $v = 0$ should bid $B(0) = 0$, which means $R(0) = 0$. As v approaches 0, $\frac{C}{v^n}$ would approach infinity unless $C = 0$. Thus, $C = 0$ and:

$$R(v) = \frac{n-1}{n} \tag{10}$$

Therefore, $B(v) = vR(v) = \frac{n-1}{n}v$, which is the symmetric BNE bidding strategy for the first-price auction. Letting $n = 3$ for our results gives us the result.

A.2.3 Third-price sealed-bid auction (TPSB) risk-neutral BNE

In a third-price sealed-bid auction with $n \geq 3$ risk-neutral bidders whose valuations are independently drawn from a uniform distribution on $[0, \bar{v}]$, the symmetric Bayesian Nash equilibrium bidding strategy is $B(v) = \frac{n-1}{n-2}v$.

Proof. In a third-price sealed-bid auction, the winner pays the third-highest bid. We use Myerson (1981)'s revenue equivalence result to derive the equilibrium.

First, note that the expected payment by a bidder with valuation v in a second-price auction is:

$$P_2(v) = E[\textit{second-highest bid} \mid v \text{ is highest}] \cdot Pr(v \text{ is highest}) \tag{11}$$

For uniform distributions on $[0, \bar{v}]$, the expected value of the second-highest of $n-1$ values (given that they are all below v), is $\frac{n-2}{n}v$. And the probability that v is the highest is $\left(\frac{v}{\bar{v}}\right)^{n-1}$.

By revenue equivalence and given risk-neutrality, the expected payment in any efficient auction mechanism must equal this amount. In a third-price auction, the expected payment is:

$$P_3(v) = E[\textit{third-highest bid} \mid v \text{ is highest}] \cdot Pr(v \text{ is highest}) \tag{12}$$

The expected value of the third-highest of $n-1$ values, given that they are all below v, is $\frac{n-3}{n-1}v$.

Let $B_3(v)$ be the equilibrium bidding function for the third-price auction. For revenue equivalence to hold:

$$B_3(v) \cdot \left(\frac{v}{\bar{v}}\right)^{n-1} = \frac{n-2}{n}v \cdot \left(\frac{v}{\bar{v}}\right)^{n-1} \tag{13}$$

Solving for $B_3(v)$:

$$B_3(v) = \frac{n-1}{n-2}v \tag{14}$$

And we have the symmetric BNE bidding strategy for the third-price auction. This proof method is classic in the literature and only uses the 2nd-order statistic with revenue equivalence. Once again, giving $n = 3$ gives us the result.

A.2.4 All-pay auction risk-neutral BNE

In an all-pay auction with n risk-neutral bidders whose valuations are independently drawn from a uniform distribution on $[0, \bar{v}]$, the symmetric Bayesian Nash equilibrium bidding strategy is $B(v) = \frac{n-1}{n} \cdot \frac{v^n}{\bar{v}^{n-1}}$.

Proof. In an all-pay auction, all bidders pay their bids regardless of whether they win or lose. Consider a bidder with valuation v who bids b. If all other bidders follow a strictly increasing bidding strategy $B(v)$, the probability of winning is $[F(B^{-1}(b))]^{n-1}$, where F is the distribution function of valuations.

The expected payoff is:

$$\pi(b, v) = v \cdot [F(B^{-1}(b))]^{n-1} - b \tag{15}$$

Taking the first-order condition with respect to b:

$$\frac{\partial \pi}{\partial b} = v \cdot (n-1) \cdot [F(B^{-1}(b))]^{n-2} \cdot f(B^{-1}(b)) \cdot \frac{d}{db} B^{-1}(b) - 1 = 0 \tag{16}$$

In a symmetric equilibrium, $B^{-1}(b) = v$, so $\frac{d}{db} B^{-1}(b) = \frac{1}{B'(v)}$. For a uniform distribution, $F(v) = \frac{v}{\bar{v}}$ and $f(v) = \frac{1}{\bar{v}}$. Substituting these into the FOC:

$$v \cdot (n-1) \cdot \left(\frac{v}{\bar{v}}\right)^{n-2} \cdot \frac{1}{\bar{v}} \cdot \frac{1}{B'(v)} - 1 = 0 \tag{17}$$

Solving for $B'(v)$:

$$B'(v) = \frac{v \cdot (n-1) \cdot v^{n-2}}{\bar{v}^{n-1}} = \frac{(n-1) \cdot v^{n-1}}{\bar{v}^{n-1}} \tag{18}$$

Integrating both sides:

$$B(v) = \int_0^v \frac{(n-1) \cdot x^{n-1}}{\bar{v}^{n-1}} dx = \frac{n-1}{n} \cdot \frac{v^n}{\bar{v}^{n-1}} + C \tag{19}$$

where C is a constant of integration. Since a bidder with valuation $v = 0$ would never want to bid a positive amount, $B(0) = 0$, which implies $C = 0$. Therefore, the equilibrium bidding strategy is:

$$B(v) = \frac{n-1}{n} \cdot \frac{v^n}{\bar{v}^{n-1}} \tag{20}$$

This is the symmetric Bayesian Nash equilibrium for the all-pay auction with risk-neutral bidders. Note that unlike the FPSB auction, the bid is not linear in the valuation, and for high valuations, the bid can represent a significant fraction of the valuation, approaching $\frac{n-1}{n} v$ as v approaches $\bar{v}$.

For the uniform distribution on $[0, \bar{v}]$, we have specifically derived the equilibrium bidding functions:

$$B_{FPSB}(v) = \frac{n-1}{n} v \tag{21}$$

$$B_{SPSB}(v) = v \tag{22}$$

$$B_{TPSB}(v) = \frac{n-1}{n-2} v \tag{23}$$

$$B_{All-pay}(v) = \frac{n-1}{n} \cdot \frac{v^n}{\bar{v}^{n-1}} \tag{24}$$

Letting $n = 3$ and $\bar{v} = 99$ gives us the result.

A.3 Risk-averse equilibria

Below, we prove the form of CARA and CRRA equilibria of the FPSB, TPSB and all-pay auctions. For all proofs, we assume symmetric and monotone bidding functions, and we fix the setting to that of $n = 3$ buyers drawing independent, private values uniformly from $[0, 99]$. The SPSB equilibrium is in dominant strategies, so it holds irrespective of risk-attitudes and requires no new analysis.

A.3.1 FPSB, CARA

First, the FPSB auction with CARA Utility. Consider a first-price sealed-bid auction with $n = 3$ risk-averse bidders, each having CARA utility

$$U(x) = 1 - e^{-\lambda x},$$

with $\lambda > 0$. Then there exists a unique symmetric Bayesian-Nash equilibrium with bidding function

$$B_{\text{CARA}}(v) = v - \frac{1}{\lambda} \ln \left[3 - 2e^{-\lambda v} \right].$$

Proof. If a bidder with value v deviates to bid b, her expected utility is

$$\text{EU}(b; v) = \left(\frac{B^{-1}(b)}{99} \right)^2 [1 - e^{-\lambda(v-b)}] + \left(1 - \left(\frac{B^{-1}(b)}{99} \right)^2 \right) [1 - e^{-\lambda \cdot 0}],$$

where $B^{-1}(b)$ is the value that would lead to bid b under the equilibrium strategy. Since $e^{-\lambda \cdot 0} = 1$, the second term simplifies to zero, leaving

$$\text{EU}(b; v) = \left(\frac{B^{-1}(b)}{99} \right)^2 [1 - e^{-\lambda(v-b)}].$$

In equilibrium, we set $b = B(v)$, so $B^{-1}(b) = v$, and the expected utility becomes

$$\text{EU}(B(v); v) = \left(\frac{v}{99} \right)^2 [1 - e^{-\lambda(v - B(v))}].$$

The first-order condition requires that the derivative of this expression with respect to $B(v)$ equals zero. Following the approach in Riley and Samuelson (1981), we derive a differential equation of the form:

$$\frac{d}{dv} \left[\left(\frac{v}{99} \right)^2 \cdot e^{-\lambda(v - B(v))} \right] = 0$$

Through integration and applying the boundary condition $B(0) = 0$, we obtain the solution:

$$B_{\text{CARA}}(v) = v - \frac{1}{\lambda} \ln \left[3 - 2e^{-\lambda v}\right].$$

To verify that this is indeed the correct solution, we can check the boundary condition:

$$B_{\text{CARA}}(0) = 0 - \frac{1}{\lambda} \ln[3 - 2e^0] = 0 - \frac{1}{\lambda} \ln[3 - 2] = 0 - \frac{1}{\lambda} \ln[1] = 0.$$

As $\lambda \to 0$ (i.e., risk neutrality), $B_{\text{CARA}}(v) \to \frac{2}{3}v$, the risk-neutral equilibrium derived above.

A.3.2 FPSB, CRRA

Next, the FPSB auction with CRRA Utility. Consider a first-price sealed-bid auction with $n = 3$ risk-averse bidders, each having CRRA utility

$$U(x) = \begin{cases} \frac{x^{1-\rho}}{1-\rho}, & \rho \neq 1, \\ \ln(x), & \rho = 1, \end{cases}$$

and value $v \in [0, 99]$. Then there exists a unique symmetric Bayesian-Nash equilibrium where the bidding function is

$$B_{\text{CRRA}}(v) = \frac{2}{2+\rho} \cdot v.$$

Proof. We begin by formulating the expected utility for a bidder who deviates from the equilibrium strategy. Suppose other bidders use a strictly increasing bid function $\beta(\cdot)$. If a bidder with true value v deviates by bidding b, her probability of winning is

$$P_{\text{win}}(b) = \left(\frac{\beta^{-1}(b)}{99}\right)^2,$$

where β^{-1} is the inverse of the bid function. If she wins, her monetary payoff is $v - b$. Hence, her expected utility is

$$\text{EU}(b; v) = \left(\frac{\beta^{-1}(b)}{99}\right)^2 U(v - b) + \left(1 - \left(\frac{\beta^{-1}(b)}{99}\right)^2\right) U(0).$$

In a symmetric equilibrium, if the bidder's value is v, she must find it optimal to bid $b = \beta(v)$. Thus, $\beta^{-1}(b) = \beta^{-1}(\beta(v)) = v$, and her expected utility at the equilibrium strategy is

$$\text{EU}(\beta(v); v) = \left(\frac{v}{99}\right)^2 U(v - \beta(v)) + \left(1 - \left(\frac{v}{99}\right)^2\right) U(0).$$

We conjecture a linear bidding function of the form $\beta(v) = cv$ for some constant $c \in (0, 1)$. This form is mathematically tractable and consistent with results in similar auction settings. Substituting into the expected utility function:

$$\text{EU}(cv; v) = \left(\frac{v}{99}\right)^2 U(v - cv) + \left(1 - \left(\frac{v}{99}\right)^2\right) U(0).$$

For CRRA utility, we can derive the first-order condition by differentiating with respect to c. Following the method in Cox, Smith, and Walker (1988), we find that the optimal bidding coefficient is:

$$c = \frac{n-1}{n-1+\rho}.$$

With $n = 3$, we obtain $c = \frac{2}{2+\rho}$. Therefore, the unique symmetric equilibrium bidding strategy is

$$B_{\text{CRRA}}(v) = \frac{2}{2+\rho} \cdot v.$$

A.3.3 TPSB, CARA

Next, the TPSB auction with CARA Utility.

Consider a third-price sealed-bid auction with $n = 3$ risk-averse bidders, each having CARA utility

$$U(x) = 1 - e^{-\lambda x},$$

with $\lambda > 0$. Then there exists a unique symmetric Bayesian-Nash equilibrium where the bidding function is

$$B_{\text{CARA}}(v) = v + \frac{1}{\lambda} \ln(1 + \lambda v).$$

Proof. In a third-price auction with $n = 3$ bidders, a bidder with value v wins if and only if her bid exceeds the bids of all other bidders, which occurs with probability $\left(\frac{v}{99}\right)^2$ in a symmetric equilibrium.

When the bidder wins, she pays the minimum of all bids. Given that the other two bidders' values, denoted v_2 and v_3, are both less than v, the minimum of their bids is $b^*(\min\{v_2, v_3\})$. The conditional distribution of $\min\{v_2, v_3\}$ given that both are less than v has probability density function

$$f(w) = 2\frac{v-w}{v^2}, \quad 0 \le w \le v.$$

Therefore, the bidder's expected utility with CARA preferences is

$$\text{EU}(v; b^*) = \left(\frac{v}{99}\right)^2 \int_0^v [1 - e^{-\lambda(v - b^*(w))}] \cdot 2\frac{v-w}{v^2}\, dw.$$

In a symmetric equilibrium, the bidder must find it optimal to follow the strategy b^*. This yields a differential equation for $b^*(v)$. Following the approach in Kagel and Levin (1993), the differential equation can be solved explicitly for CARA utility in a third-price auction.

The resulting equilibrium bidding function is

$$b^*(v) = v + \frac{1}{\lambda} \ln(1 + \lambda v).$$

This solution satisfies the boundary condition $b^*(0) = 0$, as $b^*(0) = 0 + \frac{1}{\lambda} \ln(1 + \lambda \cdot 0) = 0$.

As $\lambda \to 0$ (approaching risk neutrality), we can use L'Hôpital's rule to find

$$\lim_{\lambda \to 0} B_{\text{CARA}}(v) = \lim_{\lambda \to 0} \left(v + \frac{1}{\lambda} \ln(1 + \lambda v) \right) = v + \lim_{\lambda \to 0} \frac{\ln(1 + \lambda v)}{\lambda}.$$

Applying L'Hôpital's rule:

$$\lim_{\lambda \to 0} \frac{\ln(1+\lambda v)}{\lambda} = \lim_{\lambda \to 0} \frac{v}{1+\lambda v} = v.$$

Therefore, $\lim_{\lambda \to 0} B_{\text{CARA}}(v) = v + v = 2v$, which is the risk-neutral equilibrium bidding strategy in a third-price auction with $n = 3$ bidders.

A.3.4 TPSB, CRRA

Next, the TPSB auction with CRRA utility.

For a third-price auction with $n = 3$ risk-averse bidders having CRRA utility $U(x) = \frac{x^{1-\rho}}{1-\rho}$ (or $\ln(x)$ if $\rho = 1$), the symmetric equilibrium bidding function $b^*(v)$ must satisfy the following ordinary differential equation:

$$b'^*(v) = \Phi(v, b^*(v), \rho)$$

with boundary condition $b^*(0) = 0$, where Φ is derived from the expected utility maximization problem.

Proof. As in the CARA case, a bidder with value v wins with probability $\left(\frac{v}{99}\right)^2$ in a symmetric equilibrium. The expected utility with CRRA preferences is

$$\text{EU}(v; b^*) = \left(\frac{v}{99}\right)^2 \int_0^v U(v - b^*(w)) \cdot 2\frac{v-w}{v^2}\, dw,$$

where $U(x) = \frac{x^{1-\rho}}{1-\rho}$ for $\rho \neq 1$, or $U(x) = \ln(x)$ for $\rho = 1$.

In equilibrium, the bidder must find it optimal to follow the strategy b^*. This yields a functional equation that characterizes the equilibrium bidding function.

For CRRA utility, the resulting differential equation does not admit a closed-form solution. We can express it as

$$b'^*(v) = \Phi(v, b^*(v), \rho),$$

which must be solved numerically for each value of $\rho > 0$, subject to the boundary condition $b^*(0) = 0$.

A.3.5 All-Pay, CARA

Consider an all-pay auction with $n = 3$ risk-averse bidders, each having CARA utility

$$U(x) = -e^{-\lambda x},$$

with $\lambda > 0$. The symmetric equilibrium bidding function $B(v)$ is characterized by the differential equation

$$2\frac{v}{99^2}[e^{-\lambda(v-B(v))} - e^{\lambda B(v)}] + \lambda \left(\frac{v}{99}\right)^2 e^{-\lambda(v-B(v))} \cdot \Psi(v, B(v), \lambda) = 0,$$

with boundary condition $B(0) = 0$, where Ψ is a function derived from the expected utility maximization problem.

Proof. In an all-pay auction, a bidder with value v who bids b has monetary payoff

$$\pi = \begin{cases} v - b, & \text{if win,} \\ -b, & \text{if lose.} \end{cases}$$

The bidder wins with probability $\left(\frac{V(b)}{99}\right)^2$, where $V(\cdot)$ is the inverse of the equilibrium bidding function $B(\cdot)$.

With CARA utility $U(x) = -e^{-\lambda x}$, the expected utility is

$$\text{EU}(b; v) = \left(\frac{V(b)}{99}\right)^2 [-e^{-\lambda(v-b)}] + \left[1 - \left(\frac{V(b)}{99}\right)^2\right] [-e^{-\lambda(-b)}].$$

Since $e^{-\lambda(-b)} = e^{\lambda b}$, this simplifies to

$$\text{EU}(b; v) = -\left(\frac{V(b)}{99}\right)^2 e^{-\lambda(v-b)} - \left[1 - \left(\frac{V(b)}{99}\right)^2\right] e^{\lambda b}.$$

In equilibrium, the bidder must find it optimal to bid $b = B(v)$. This requires that the derivative of $\text{EU}(b; v)$ with respect to b, evaluated at $b = B(v)$, equals zero.

Taking the derivative and using the chain rule with $V'(b) = \frac{1}{B'(V(b))}$, we obtain

$$2\frac{v}{99^2}[e^{-\lambda(v-B(v))} - e^{\lambda B(v)}] + \lambda\left(\frac{v}{99}\right)^2 e^{-\lambda(v-B(v))} \cdot \Psi(v, B(v), \lambda) = 0,$$

where Ψ includes terms from the derivatives of the exponential functions.

This differential equation, together with the boundary condition $B(0) = 0$, characterizes the equilibrium bidding function for the CARA case. No closed-form solution is known, so it must be solved numerically or via perturbation methods for small values of λ.

A.3.6 All-Pay, CRRA

Consider an all-pay auction with $n = 3$ risk-averse bidders, each having CRRA utility

$$U(x) = \frac{x^{1-\rho}}{1-\rho},$$

with $\rho > 0, \rho \neq 1$. The symmetric equilibrium bidding function $B(v)$ is characterized by the differential equation

$$2\frac{v}{99^2}[(v - B(v))^{1-\rho} - (-B(v))^{1-\rho}] + \Omega(v, B(v), \rho) = 0,$$

with boundary condition $B(0) = 0$, where Ω includes terms from derivatives of the power functions.

Proof. In an all-pay auction, a bidder with value v who bids b has monetary payoff

$$\pi = \begin{cases} v - b, & \text{if win,} \\ -b, & \text{if lose.} \end{cases}$$

With CRRA utility $U(x) = \frac{x^{1-\rho}}{1-\rho}$, the expected utility is

$$\text{EU}(b; v) = \left(\frac{V(b)}{99}\right)^2 \frac{(v-b)^{1-\rho}}{1-\rho} + \left[1 - \left(\frac{V(b)}{99}\right)^2\right] \frac{(-b)^{1-\rho}}{1-\rho},$$

assuming that $(-b)^{1-\rho}$ is well-defined for $b \geq 0$. In practice, we might need to add a positive constant to all payoffs to ensure they remain positive.

In equilibrium, the bidder must find it optimal to bid $b = B(v)$. This yields the first-order condition that the derivative of $\text{EU}(b; v)$ with respect to b, evaluated at $b = B(v)$, equals zero.

Differentiating with respect to b and applying the chain rule, we obtain a differential equation of the form

$$2\frac{v}{99^2}[(v - B(v))^{1-\rho} - (-B(v))^{1-\rho}] + \Omega(v, B(v), \rho) = 0,$$

where Ω includes terms from the derivatives of the power functions.

This differential equation, together with the boundary condition $B(0) = 0$, characterizes the equilibrium bidding function for the CRRA case. No closed-form solution is known, so it must be solved numerically or via perturbation methods for small values of ρ.

If $\rho > 1$ or we wish to avoid negative arguments in the CRRA utility function, we can add a positive baseline M to payoffs so that the new payoff is $(M + v - b)$ if the bidder wins and $(M - b)$ if she loses, with M large enough so that $M - b > 0$. This changes the specific form of the differential equation but preserves the general approach.

B Implementation details

B.1 Simulation process in sealed-bid auctions

During the auction, LLM agents will be first informed about the scenario. After explaining to them the rules and incentives of the aucion, they will be asked to write down a plan for the bidding.

```
You are Bidder Andy.
You are bidding with Bidder Betty, Bidder Charles.

{RULE EXPLANATION} + {INSTRUCTIONS}

Write your plans for what bidding strategies to test next.
Be detailed and precise but keep things succinct
and don't repeat yourself.
Your plan should be within 100 words.
```

Here the RULE EXPLANATIONS are specific for each type of auctions. The INSTRUCTIONS are included to incentivize the LLM agents and motivate strategic behaviors; they are universal in all the simulations we ran:

```
    Your TOP PRIORITY is to place bids
    which maximize the user's profit in the long run.
    To do this, you should explore many
    different bidding strategies,
    including possibly risky or aggressive
    options for data-gathering purposes.
    Learn from the history of previous
    rounds in order to maximize your total profit.
    Don't forget the values are redrawn
independently each round.
```

The following is the output from Bidder Andy for the first round in a FPSB auction:

> *"I plan to initially explore a strategy where I bid half of my value, then varying my bids between 30-70% of my value in subsequent rounds. Sequentially, I will introduce risk by bidding within +/- 5% of the highest bid from previous round, gathering data about other bidder's tendencies. Lastly, based on gathered data, I'll deploy a more aggressive high bid if losses amass, or low-risk bid if leading, to secure profit. Learning from each round is vital to adapt and refine strategy for profit maximization."*

Once they make the bidding plan, the system will inform the LLM agents their value towards the item in this round and ask them to make a bid according to their plan.

```
    You are Bidder Andy.
    You are bidding with Bidder Betty, Bidder Charles.

    {RULE EXPLANATION} + {INSTRUCTIONS}

    Your value is 73. Your plan is {PLAN}
    FOLLOW YOUR PLAN.
    How much would you like to bid?
    Give your response with a single number and no other texts,
    e.g. 1, 44
```

In this example, Bidder Andy valued the item at $73. And according to his plan, he only wanted to bid 30-70% of his value. So, he bid *"36.5"* in this round.

After receiving the bids from all the LLM bidder, the system determines the winner and the profits for each bidders. All the bids and the winner's profit are returned back to the bidders along with whether they win the auction in this round. The part of information is included in the HISTORY variable. For sealed-bid auctions, the HISTORY includes the all the bids, the clearing price, winner's profit, this agent's value, agent's bid and whether they won the auction in that round. An example of the history briefing is provided below:

> *"The previous round histories along with your plans are: In round 0, Your value was 86, you bid 86.0, and your profit was 0.0. Your total profit is 0.0. All the bids for this round were 86.0, 56.0, 21.0. The highest bidder won with a bid of 86.0 and would've preferred to bid 57.0. The winner's profit was 0.0. Did you win the auction: Yes. "*

Before entering the next round, the bidders are asked to do a reflection upon their bidding strategy and the result.

```
You are Bidder Andy.
You are bidding with Bidder Betty, Bidder Charles.

{RULE EXPLANATION} + {INSTRUCTION}

The previous round history is: {HISTORY}

Do a counterfactual analysis of the last round.
REMEMBER that your goal is to win the bid and make higher profits.
REMEMBER YOUR PAYMENT IS YOUR BID IF YOU WIN.
Let's think step by step.
Start your reflection with
'If I bid down by .., I could...
If I bid up by ..., I could...'
LIMIT your OUTPUT within 100 words.
```

An example of Bidder Andy's reflection is provided below. Since he didn't win the auction and incurred a zero profit, his reflection was that bidding up closer to his true value would help him win the game and secure a net profit.

> *"If I bid down by a few dollars, say around $30, I likely still wouldn't have won given the winning bid was $57. I would've merely saved my bid money. If I bid up by $21, matching the winning bid of $57, I may or may not have won due to the random resolution of ties. However, if I'd won, my profit would have been $16, realizing the difference between my value of $73 and my raised bid. So, bidding aggressively could potentially increase my profits in such scenarios. "*

In the following new round, all the bidders are asked to update their plans according to the history of the last round and their reflections.

```
You are Bidder Andy.
You are bidding with Bidder Betty, Bidder Charles.

{RULE EXPLANATION} + {INSTRUCTION}

The previous round history is: {HISTORY}
After careful reflection on previous bidding,
your analysis for last round is {REFLECTION}.

Learn from your previous rounds,
Let's think step by step to make sure we make a good choice.
Write your plans for what bidding strategies to test next.
Be detailed and precise but keep things succinct and don't repeat yourself.
LIMIT your plan to 50 words.
```

Here, we show the updated plan for the second round of the same Bidder Andy. After learning from the lessons of bidding only 50% of the value and losing the game, this time the agent updates their plan to the following:

> *"I plan to test higher bidding strategies, around two-thirds of the prize value to foster winning chances. Additionally, if my prize value is extremely high, I will bid aggressively for data-gathering."*

This plan-bid-reflection loop will iterate until the end of the auction. After the last bid, no reflection will be queried. In the IPV setting, there are overall 15 rounds for all the sealed-bid auctions. In the affiliated and common value auctions, there are overall 10 rounds.

B.2 Simulation process in Ascending Clock auctions

In Ascending Clock Auctions, each round consists of multiple clock cycles, during which every bidder is asked whether they want to stay or drop out at the current clock price. In the first round, bidders are reminded of the auction rules, similar to those outlined in Section B.1. The auction starts with an initial price of 0, which increases incrementally until only one bidder remains or two bidders drop out simultaneously, in which case the winner is chosen randomly. The detailed prompt is listed as follows, with variables enclosed in brackets:

```
You are Bidder Andy.
You are bidding with Bidder Betty, Bidder Charles.

{RULE EXPLANATION} + {INSTRUCTIONS}

Your value towards to the prize is 26 in this round.
The current price in this clock cycle is {current_price}.
The price for next clock cycle is {current_price + increment}.

The previous bidding history is: {transcript}.

{Asking Prompt}
```

To implement plan-bid-reflection within each clock cycle, we modified the Asking Prompt as follows, inspired by Fish et al. (2024):

```
Do you want to stay in the bidding?
If you choose yes, you can keep bidding for next clock.
If you choose No, you will exit and
have no chance to re-enter the bidding.
Your response must use these EXACT
tags below. You must output the ACTION.
```
```
<PLAN>
[Write your plans for what bidding strategies to test next.
Be detailed and precise but keep things succinct
```

```
and don't repeat yourself.
LIMIT your plan to 50 words. ] </PLAN>
<ACTION> Yes or No </ACTION>
<REFLECTION> think about how to justify your choice </REFLECTION>
```

In AC, if LLM bidders didn't decide to drop out in the first round, they are shown the previous bidding history for the next clock cycle. The history includes the price and how many people had dropped out. And the transcript variable will be the following format:

```
The previous biddings are:
['In clock round 1, the price was 1,
no players dropped out'].
```

In AC-B, we don't show the previous biddings and in each clock cycle, we directly query LLM agents' decision. So the transcript variable will be None.

A clock auction round ends when only one bidder remains, who is declared the winner, or when two bidders drop out simultaneously, in which case the winner is chosen randomly. Each auction consists of 10 independent rounds with affiliated value settings.

B.3 Simulation process in eBay auctions

For eBay auctions, we discretize the continuous bidding period into 10 periods. Each period, LLM agents decide whether to increase their bid or hold their current bid. The bidding process follows a structured format, where bidders act in a predefined sequence each period (e.g., Charles, then Alice, then Betty). Agents are informed of past price changes through a transcript, such as "On day 1, the price changed to 1. On day 2, the price changed to 3."

The prompt provides key details, including the bidder's private value, the total number of bidding days, the current day, the bidding order, previous bids, previous proxy bids, and the current price.

```
You are Bidder Andy.
You are bidding with Bidder Betty, Bidder Charles.

{RULE EXPLANATION} + {INSTRUCTIONS}

Your private value for this item is ${private_value}.
This is the maximum amount you are
willing to pay. Keep this value private.
There are in total {total_periods} days
of bidding and this is day {current_period}.
{ordering}.

Your previous bids are {previous_bid}.
If you have already placed bids, you can only increase your bid
or hold your current bid.
The previous proxy bids are: {transcript}.
```

```
The current price is {current_price}.
{Asking Prompt}
```

The decision-making process is again guided by a structured asking prompt that enforces bidding rules and strategic reflection.

```
Your response must use these EXACT
tags below. Don't miss the amount.
'''
<PLAN>[ Write your plans for what
bidding strategies to test next.
Be detailed and precise but keep
things succinct and don't repeat yourself.
LIMIT your plan to 50 words. ] </PLAN>
<CHECK> your last bid is
{{last_bid_amount}}, you cannot bid
lower that this value </CHECK>
<ACTION> BID or HOLD  </ACTION>
<AMOUNT> if BID, enter a number here,
e.g. 1, 44. If HOLD, enter 0 </AMOUNT>
'''
```

C One-Round Sealed-Bid Auctions and the Disappearance of Overbidding

A key puzzle in our multi-round settings was: Why do some LLM bidders bid above their assigned value? From the plans and reflections provided by the LLMs, we found two main motivations. First, in multi-round games, the LLMs sometimes purposely inject random or high bids in order to learn how their opponents play or to appear "unpredictable." For example, one LLM bidder declared:

> *I'll adopt a balanced strategy, bidding 65% of my value. I'll also introduce random bids occasionally to disrupt predictability. Monitoring competitor's bids remains essential to adjust my strategy accordingly.*

Second, overbidding in Second Price Sealed-Bid (SPSB) auctions is a well-known "failure mode" even among humans. Although bidding one's true value is a dominant strategy in SPSB, substantial experimental evidence finds that people systematically overbid. Kagel and Levin (1993) report that 67.2% of participants overbid, and even experienced bidders studied by Garratt, Walker, and Wooders (2012) overbid 37.5% of the time. In fact, the persistent puzzle of overbidding in SPSB helped motivate subsequent theoretical work (e.g., Shengwu Li's paper on obviously strategy-proof mechanisms).

To test whether multi-round considerations (information gathering, unpredictability, or learning) were driving the overbidding phenomenon in our LLM bidders, we re-ran all sealed-bid auctions in a strictly single-round setting. We removed any mention that the game might continue beyond one period, and we explicitly stated:

```
Your top priority is to place a bid
that maximizes your expected profit.
```

Fig 8 shows the resulting one-round bids across four private-value sealed-bid auction formats. Remarkably, the overbidding phenomenon essentially disappears in one-round interactions.

For First-Price Auction, the bids lie below the diagonal line for most values, much like the Bayesian Nash equilibrium prediction. For SPSB, bids track closely with true values—aligning better with the truthful dominant strategy than in multi-round sessions. For TPSB, the LLM bidders never bid above their value; many actually shade their bids down quite aggressively, often below the theoretical equilibrium curve. For All-Pay Auction, overbidding is also absent here, though some moderate upward bidding at higher valuations remains.

Taken together, these results indicate that LLMs' attempts to be strategic over over repeated interactions is a key factor for overbidding observed in our experiments. When presented with only one shot and explicitly instructed to maximize profit, LLM bidders' tendency to bid above their assigned value virtually vanishes.

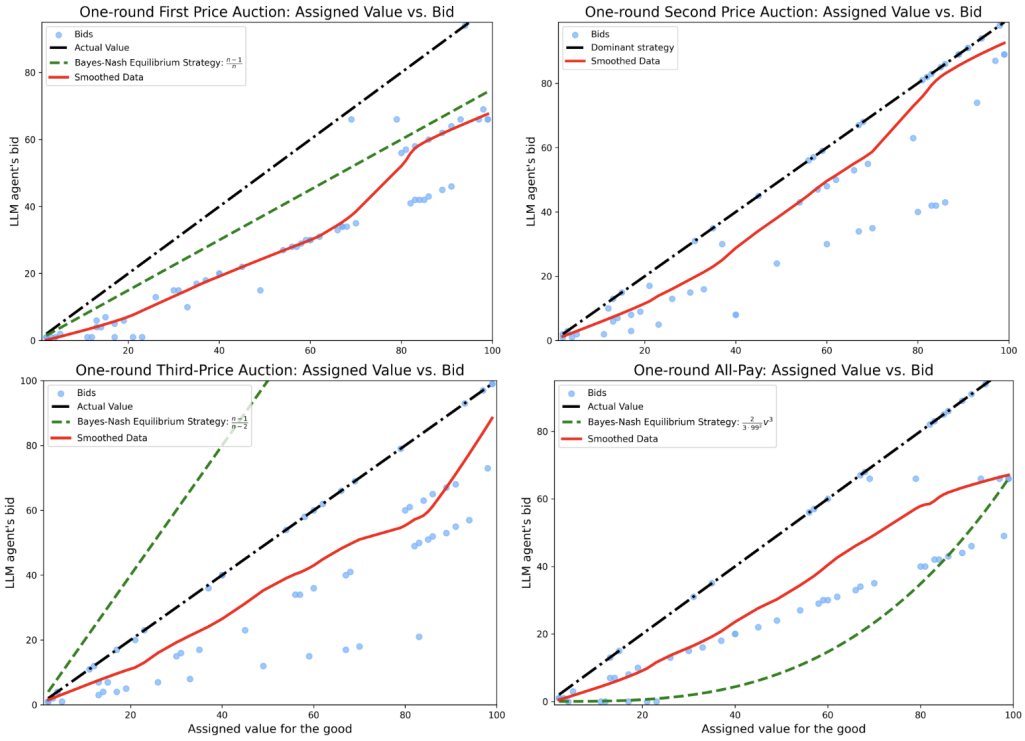


Figure 8: One-round Sealed-bid auction under IPV

D Robustness Checks

D.1 Varying the Number of agent in First-price sealed-bid auctions

In FPSB auctions, the theoretically optimal bidding strategy varies based on the number of agents participating. As the number of agents increases, we would expect individual bidders to reduce the extent of bid shading.

To examine this effect, we conducted experiments with FPSB auctions involving 4 and 5 agents, as shown in Fig 9. As the number of agents increases, the Loess-smoothed data curve remains consistently higher than, or approximately aligned with, the predictions of the Bayes-Nash equilibrium. Moreover, we observed variations below the theoretical optimum, along with a few data points exceeding the predicted values. These findings align with existing empirical evidence for first-price auctions, suggesting that the observed patterns persist across different settings.

D.2 Varying Currency in SPSB and FPSB auctions

Additionally, we also tested the prompts in different currencies (specifically, the Euro, the Ruble, the Yen and the Rupee) with the results presented in Figure 10. As is evident from the plots for both the SPSB and FPSB auctions, the key features of the FPSB and SPSB simulation results reported in the main text remain consistent and unchanged in these variations.

D.3 Varying Language in rule explanations in SPSB and FPSB

We also tested our existing prompts using different languages (specifically, Spanish, Chinese, Russian and Hindi) with the results presented in Figure 11. It is interesting that the Russian version of FPSB appears to be a stark outlier, with the LLM agent bidding close to the true value even in the first-price auction. For other languages, bids in the second-price auction are generally higher than those in the first-price auction, which is qualitatively unchanged from the results reported in the main body of paper.

E An ablation study with an 'Out-of-the-Box' LLM agent

Throughout Section 3.1, we endowed agents with the ability to plan and reflect as they played the FPSB and SPSB auctions. This 'plan-bid-reflect' loop reflects a modest instantiation of the chain-of-thought paradigm common in the machine learning literature and previous LLM auction experiments (Fish et al., 2024; Wei et al., 2022). However, one might reasonably ask how much of the sophisticated play we observe is actually due to this ability to plan and reflect.

Hence, we also report results of the experiments of Section 3.1 with 'Out-of-the-box' agents—namely, we directly query LLM agents (after providing them with auction rules and their value, as per our simulation process) *without* eliciting their plans or allowing them to reflect on the results of the previous rounds. These IPV experiment results are reported in Figure 12.

Overall, the bidding behaviors exhibited by the agents remain largely monotonic. However, interesting differences emerged between the plan-bid-reflect agents and the out-of-the-box agents. With the out-of-the-box LLM bidders there is now little difference in the bidding behavior between the First-Price and Second-Price auction, a stark difference from the plan-bid-reflect LLM bidders. Now, and unlike in human-subject experiments, no overbidding beyond the agent's value was observed in SPSB auctions. Specifically, in the left panel of Figure 12, no points lie above the bid=value line.

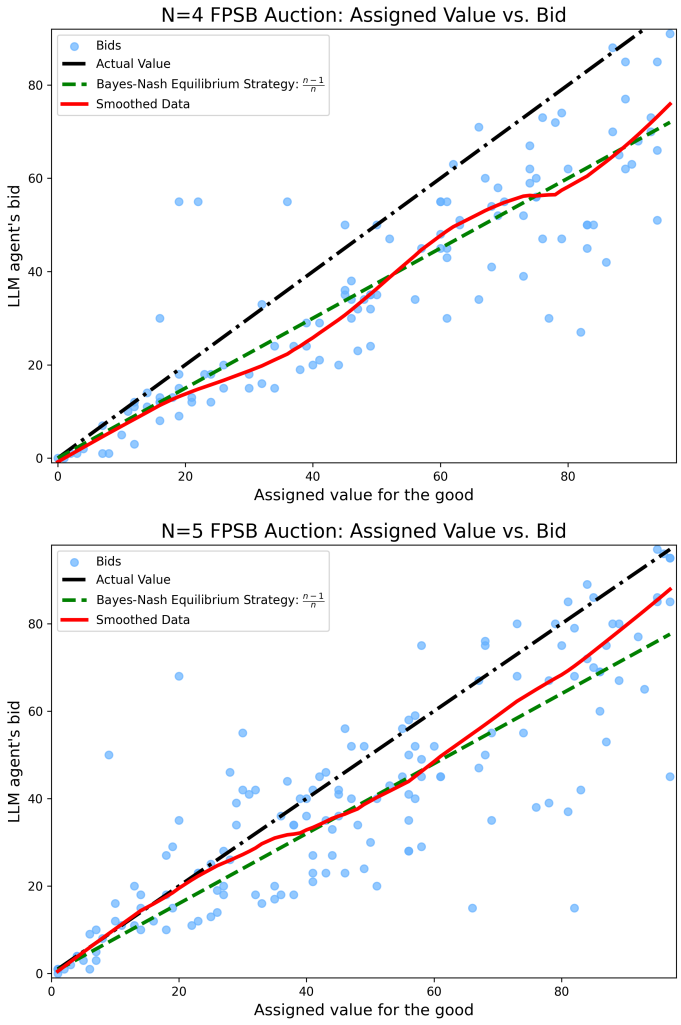


Figure 9: Robustness check of Number of agent in FPSB [DCP: IPV setting?]

F Supplementary Materials in eBay auctions

G Rules Explanation Prompts

Below are the rules and instructions of each setting for our experiments, which are argument passed to {RULE EXPLANATION} in earlier section.

For the First-Price Sealed-Bid prompt in the IPV setting, the rule is:

```
In this game, you will participate in an auction for a prize
    against {{num_bidders}} other bidders. You will play this
    game for {{n}} rounds.
```

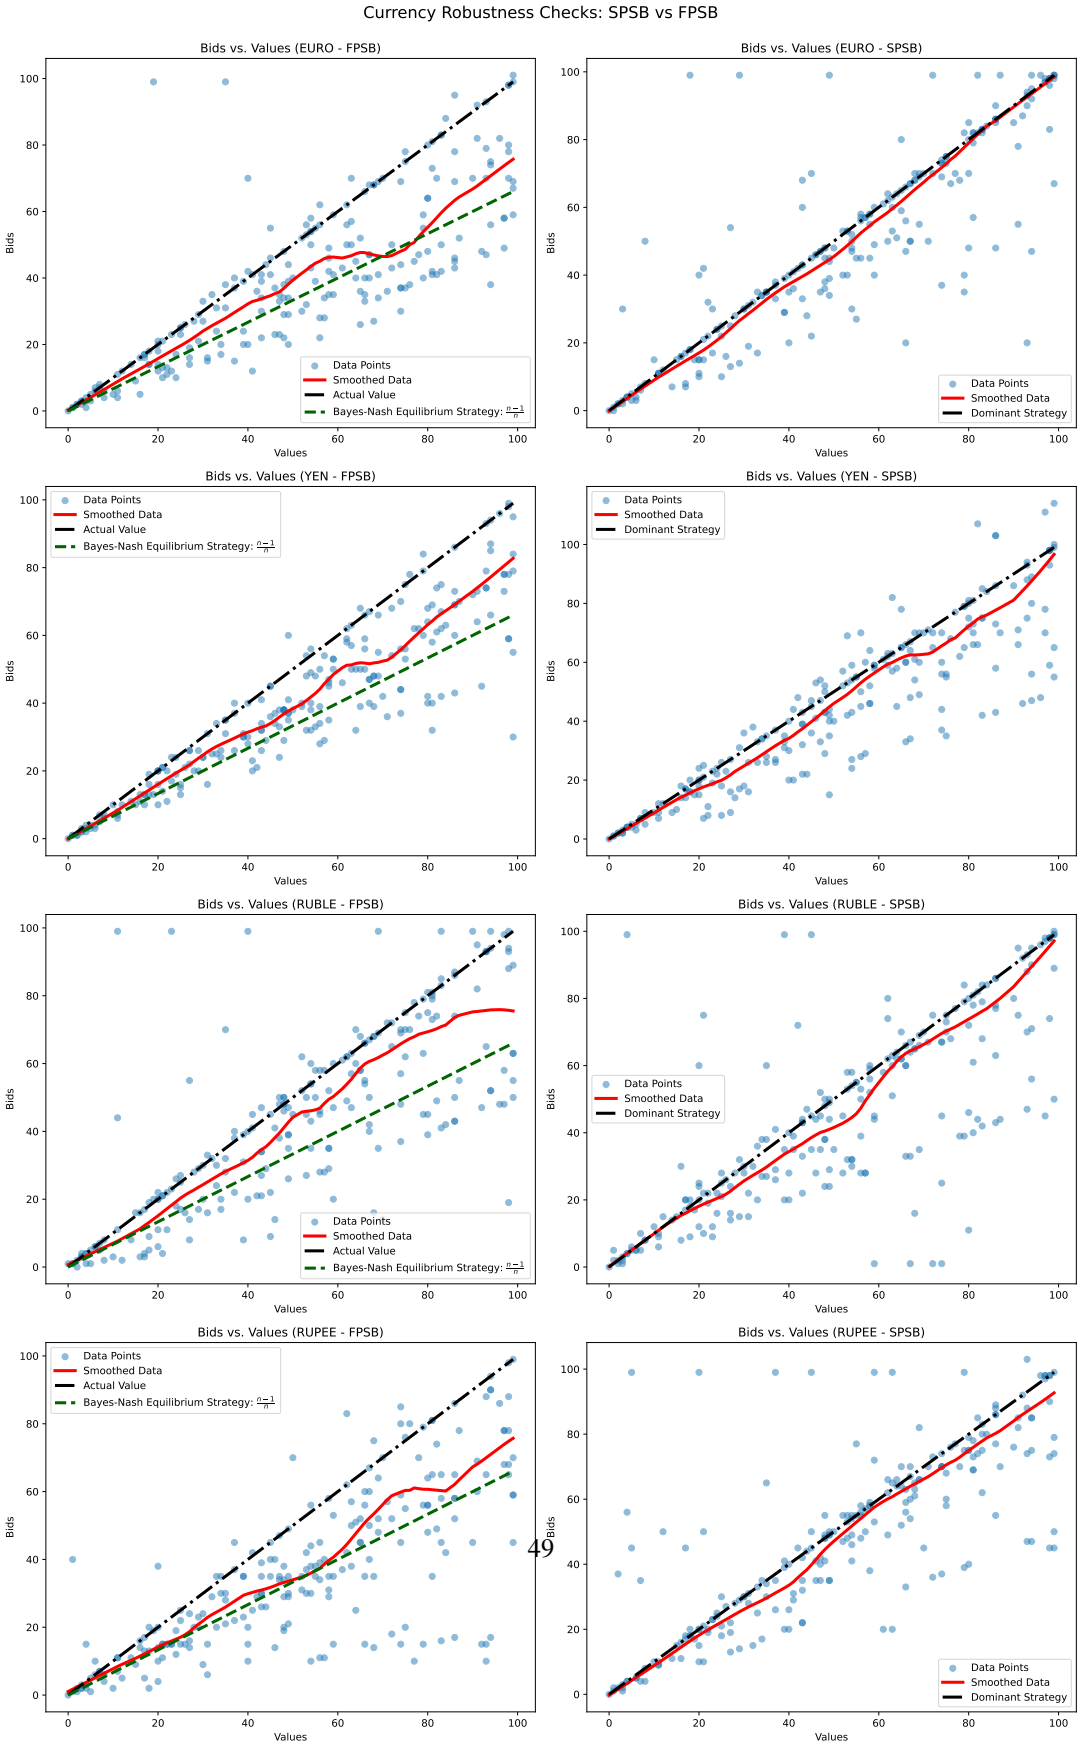


Figure 10: Robustness check with Currency Variation under IPV

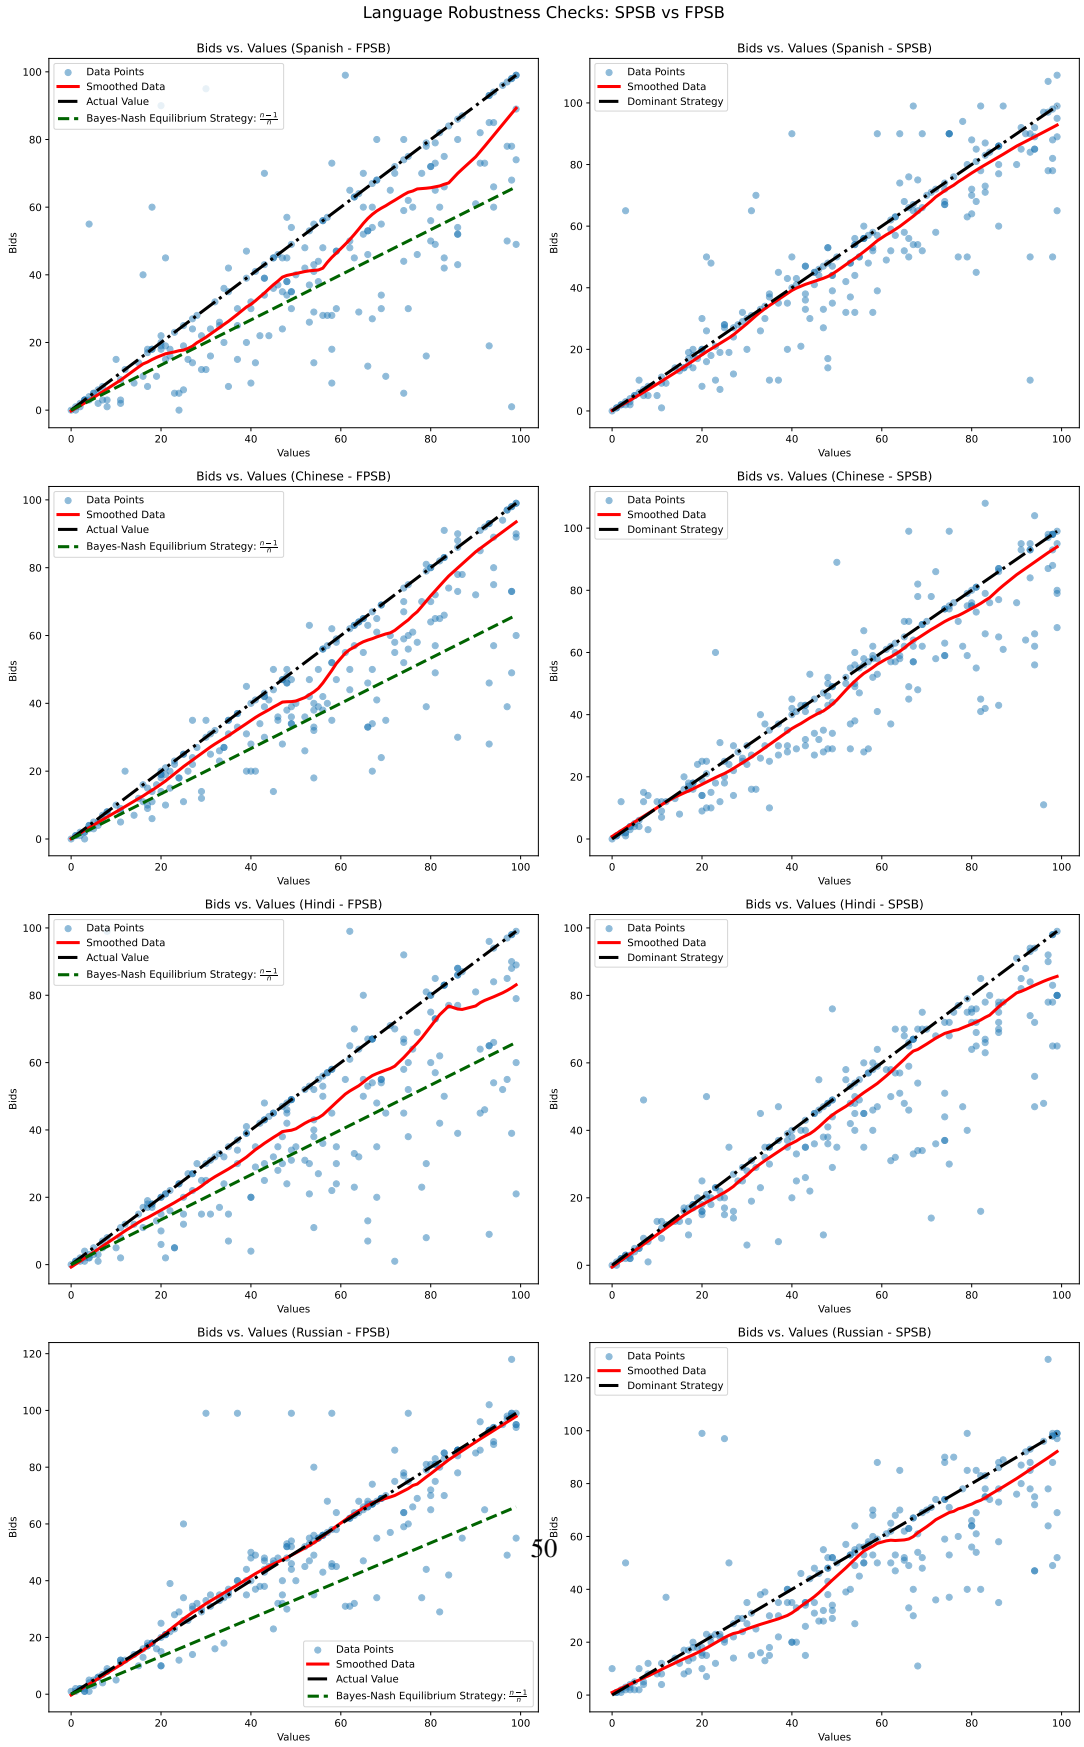


Figure 11: Robustness check with Language Variation with IPV

Auction Type 1	Auction Type 2	t-statistic	p-value
basic_proxy	closing_rule	0.020	0.984
basic_proxy	hidden_reserve_40	−0.026	0.98
basic_proxy	hidden_reserve_50	−0.096	0.923
basic_proxy	hidden_reserve_60	−0.071	0.944
basic_proxy	closing_rule_hidden_40	−0.109	0.913
basic_proxy	closing_rule_hidden_50	−0.109	0.913
basic_proxy	closing_rule_hidden_60	−0.119	0.905
closing_rule	hidden_reserve_40	−0.047	0.963
closing_rule	hidden_reserve_50	−0.120	0.905
closing_rule	hidden_reserve_60	−0.093	0.926
closing_rule	closing_rule_hidden_40	−0.133	0.895
closing_rule	closing_rule_hidden_50	−0.133	0.895
closing_rule	closing_rule_hidden_60	−0.144	0.886
hidden_reserve_40	hidden_reserve_50	−0.071	0.944
hidden_reserve_40	hidden_reserve_60	−0.045	0.964
hidden_reserve_40	closing_rule_hidden_40	−0.083	0.934
hidden_reserve_40	closing_rule_hidden_50	−0.083	0.934
hidden_reserve_40	closing_rule_hidden_60	−0.093	0.927
hidden_reserve_50	hidden_reserve_60	0.026	0.98
hidden_reserve_50	closing_rule_hidden_40	−0.013	0.99
hidden_reserve_50	closing_rule_hidden_50	−0.013	0.99
hidden_reserve_50	closing_rule_hidden_60	−0.020	0.984
hidden_reserve_60	closing_rule_hidden_40	−0.038	0.97
hidden_reserve_60	closing_rule_hidden_50	−0.038	0.97
hidden_reserve_60	closing_rule_hidden_60	−0.046	0.963
closing_rule_hidden_40	closing_rule_hidden_50	0.000	1
closing_rule_hidden_40	closing_rule_hidden_60	−0.006	0.995
closing_rule_hidden_50	closing_rule_hidden_60	−0.006	0.995

Table 10: Pairwise t-tests for seller revenue differences across auction types. None of the comparisons show significant revenue differences.

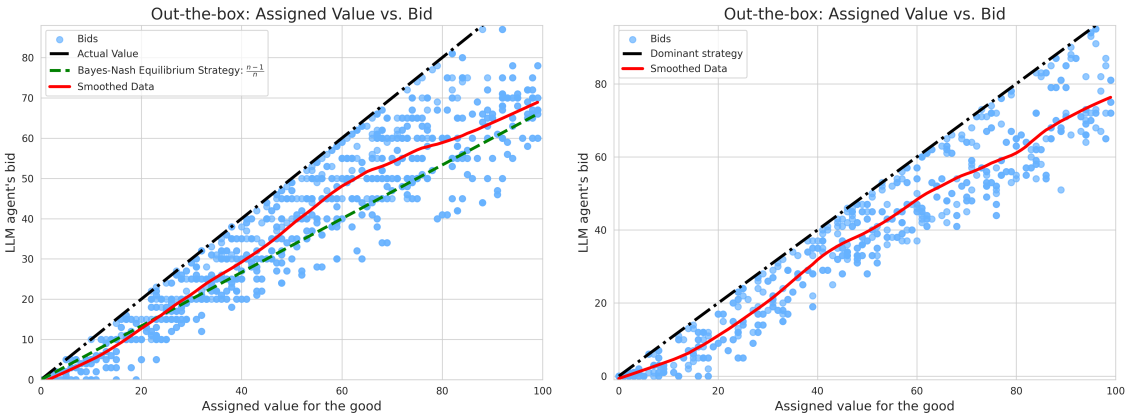


Figure 12: Out-of-the-Box LLM agent in FPSB (left) and SPSB (right) auctions under IPV setting. The blue dots represent bids, the black dotted line is where bid = value (the SPSB dominant strategy), and the green line on the right represents the BNE for FPSB auctions.

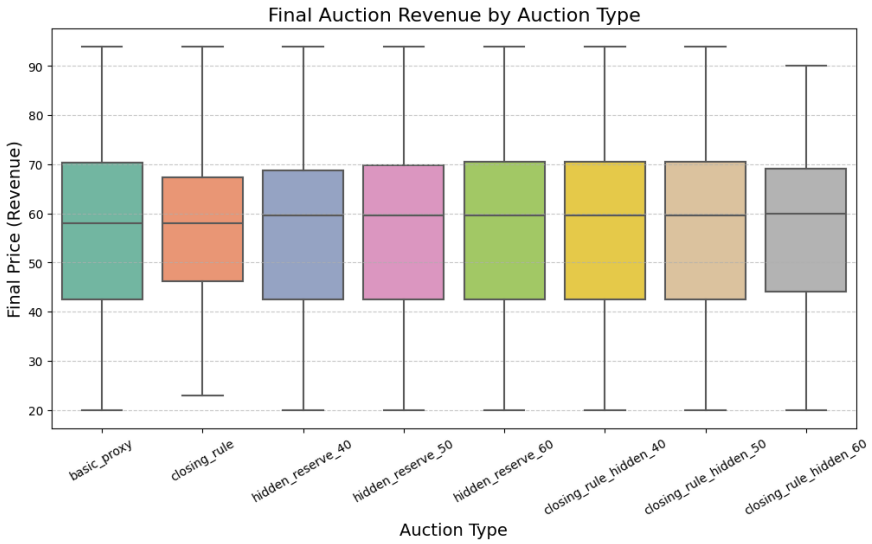


Figure 13: **eBay Revenue by Auction Type, IPV environment**. In all auction environments, design rules don't meaningfully change profits, just the speed of price discovery.

```
At the start of each round, bidders will see their value for
    the prize, randomly drawn between \$0 and \${{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. If
    you win, your earnings will increase by your value for the
    prize, and decrease by your bid. If you don't win, your
    earnings will remain unchanged.
After each auction, we will display all bids and the winner's
    profits. Ties for the highest bid will be resolved randomly.
```

For the Second-Price Sealed-Bid prompt in the IPV setting, the rule is:

```
In this game, you will participate in an auction for a prize
    against {{num_bidders}} other bidders. You will play this
    game for {{n}} rounds.
At the start of each round, bidders will see their value for
    the prize, randomly drawn between $0 and ${{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. If you win, your earnings will increase by your value
    for the prize, and decrease by the second-highest bid. If
    you don't win, your earnings will remain unchanged.
After each auction, we will display all bids and the winner's
    profits. Ties for the highest bid will be resolved randomly.
```

For the Third-Price Sealed-Bid prompt in the IPV setting, the rule is:

```
In this game, you will participate in an auction for a prize
    against {{num_bidders}} other bidders. You will play this
    game for {{n}} rounds.
At the start of each round, bidders will see their value for
    the prize, randomly drawn between $0 and ${{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 third-highest
    bid. If you win, your earnings will increase by your value
    for the prize, and decrease by the third-highest bid. If you
    don't win, your earnings will remain unchanged. If there
    are less than three bids, no one will win the auction.
After each auction, we will display all bids and the winner's
    profits. Ties for the highest bid will be resolved randomly.
```

For the All-Pay Sealed-Bid prompt in the IPV setting, the rule is:

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their value for
   the prize, randomly drawn between $0 and ${{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. All bidders (including the
   winner) pay their submitted bid. If you win, your earnings
   will increase by your value for the prize, and decrease by
   your bid. If you don't win, your earnings will still
   decrease by your bid.
After each auction, we will display all bids and all bidders'
   profits. Ties for the highest bid will be resolved randomly.
```

For the Ascending Clock auction (AC) in the APV setting, the rule is

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their value for
   the prize. Your value for the prize will be calculated as
   follows:
1. First we will randomly draw a common value between {{
   common_low}} and {{common_high}}, with all values equally
   likely.
2. Then, for each bidder, a private taste adjustment will be
   drawn between 0 and {{private}}, with all values equally
   likely.
Your value for the prize is equal to the common value plus your
    private taste adjustment. You will not learn the common
   value or your private taste adjustment separately. This
   means that each person in your group may have a different
   value for the prize. However, if 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 at 0
   and counts upwards in {{increment}} USD increments, up to a
   maximum of {{common_high + private}}. At any point, you can
   choose to leave the auction, and anytime a bidder leaves, we
    will broadcast that information to all the remaining
   bidders.
```

```
When there is only one bidder left in the auction, that bidder
   will win the prize at the current price. If you win, your
   earnings will increase by your value for the prize, and
   decrease by the current price. If you don't win, your
   earnings will remain unchanged.
After each auction, we will display all bids and the winner's
   profits. Ties for the highest bid will be resolved randomly.
```

For the Ascending Clock auction without dropping information (AC-B) in the APV setting, the rule is

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their value for
   the prize. Your value for the prize will be calculated as
   follows:
1. First we will randomly draw a common value between {{
   common_low}} and {{common_high}}, with all values equally
   likely.
2. Then, for each bidder, a private taste adjustment will be
   drawn between 0 and {{private}}, with all values equally
   likely.
Your value for the prize is equal to the common value plus your
    private taste adjustment. You will not learn the common
   value or your private taste adjustment separately. This
   means that each person in your group may have a different
   value for the prize. However, if 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 at 0
   and counts upwards in {{increment}} USD increments, up to a
   maximum of {{common_high + private}}. At any point, you can
   choose to leave the auction, but we will not tell any bidder
    when someone leaves.
When there is only one bidder left in the auction, that bidder
   will win the prize at the current price. If you win, your
   earnings will increase by your value for the prize, and
   decrease by the current price. If you don't win, your
   earnings will remain unchanged.
After each auction, we will display all bids and the winner's
   profits. Ties for the highest bid will be resolved randomly.
```

For the Second-Price Sealed-Bid Auction in the APV setting, the rule is

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their value for
   the prize. Your value for the prize will be calculated as
   follows:
1. First we will randomly draw a common value between {{
   common_low}} and {{common_high}}, with all values equally
   likely.
2. Then, for each bidder, a private taste adjustment will be
   drawn between 0 and {{private}}, with all values equally
   likely.
Your value for the prize is equal to the common value plus your
    private taste adjustment. You will not learn the common
   value or your private taste adjustment separately. This
   means that each person in your group may have a different
   value for the prize. However, if you have a high value, it
   is more likely that other people in your group have a high
   value.
After learning your value, you will submit a bid privately at
   the same time as the other bidders. Bids must be between $0
   and ${{common_high + private}} in ${{increment}} increments.
The highest bidder wins the prize and pays the second-highest
   bid. If you win, your earnings will increase by your value
   for the prize, and decrease by the second-highest bid. If
   you don't win, your earnings will remain unchanged.
After each auction, we will display all bids and the winner's
   profits. Ties for the highest bid will be resolved randomly.
```

For the Second-Price Sealed-Bid auction in the Common Value setting, the rule is

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their perceived
   value for the prize - a noisy measurement of the true value
   of the prize. Your perceived value for the prize will be
   calculated as follows:
1. For each round, a common value will be drawn between {{
   common_low}} and {{common_high}}, with all values equally
   likely to be drawn.
2. For each person, a private noisy adjustment will be drawn
   between -{{private}} and {{private}}, with all values
   equally likely to be drawn.
We will tell you your perceived value, the sum of the common
   value and the private noise adjustment. However, everyone's
```

```
   true value for the prize is equal to the shared common value
   .
After learning your perceived value, you will submit a bid
   privately at the same time as the other bidders. Bids must
   be between $0 and ${{common_high + private}} in ${{increment
   }} increments.
The highest bidder wins the prize and pays the second-highest
   bid. If you win, your earnings will increase by the true
   value for the prize, and decrease by the second-highest bid.
    If you don't win, your earnings will remain unchanged.
After each auction, we will display all bids and the winner's
   profits. Ties for the highest bid will be resolved randomly.
```

For the Standard eBay auction with proxy bidding (T1), the rule is

```
In this game, you will participate in an eBay auction for an
   item against {{num_bidders}} other bidders. This auction
   will last for {{num_rounds}} days. All dollar amounts in
   this game are in US Dollars ($).

Item Details:
Item Description: {{item_description}}
Item Condition: {{item_condition}}
Your Private Value: At the start of each round, bidders will
   see their value for the item, randomly drawn between $0 and
   ${{private}}, with all values equally likely. After learning
    your value for the item, you will submit a maximum bid.
   Bids must be between $0 and ${{private}} in ${{increment}}
   increments.

Auction Format:
This is an eBay auction. The auction starts at ${{start_price}}
    and will last for {{num_rounds}} days. eBay uses "proxy
   bidding." This means that if you wish to enter the auction,
   you should submit your maximum bid, and eBay will
   automatically bid on your behalf, up to your maximum, only
   as much as necessary to maintain your position as the
   highest bidder. Each day you will see the current price and
   have the opportunity to increase your maximum bid. If you do
    not want to increase your maximum bid, then output HOLD.

You may place bids at any point during the auction, even on the
    final day. No bidder will know if they (or anyone else) is
   the last bidder on the final day.

The highest bidder wins the prize and pays the auction price at
    the time the auctions clock expired. If you win, your
```

```
earnings will increase by your value for the prize, and
decrease by the final auction price. If you don't win, your
earnings will remain unchanged. Ties for the highest bid
will be resolved randomly.
```

For the eBay auction with a modified closing rule (T2), the rule is

```
In this game, you will participate in an eBay auction for an
   item against {{num_bidders}} other bidders. This auction
   will last for {{num_rounds}} days. All dollar amounts in
   this game are in US Dollars ($).

Item Details:
Item Description: {{item_description}}
Item Condition: {{item_condition}}
Your Private Value: At the start of each round, bidders will
   see their value for the item, randomly drawn between $0 and
   ${{private}}, with all values equally likely. After learning
    your value for the item, you will submit a maximum bid.
   Bids must be between $0 and ${{private}} in ${{increment}}
   increments.

Auction Format:
This is an eBay auction with a closing rule. The auction starts
    at ${{start_price}} and will last for {{num_rounds}} days.
   eBay uses "proxy bidding." This means that if you wish to
   enter the auction, you should submit your maximum bid, and
   eBay will automatically bid on your behalf, up to your
   maximum, only as much as necessary to maintain your position
    as the highest bidder. Each day you will see the current
   price and have the opportunity to increase your maximum bid.
    If you do not want to increase your maximum bid, then
   output HOLD.

This auction also has a closing rule. This means that if a new
   maximum bid is placed on the last day, the auction will
   automatically extend by another day. This extension will
   continue to be applied as long as new maximum bids are
   placed on the final day. No bidder will know if they (or
   anyone else) is the final bidder on the last day.

The highest bidder wins the prize and pays the auction price at
    the time the auctions clock expired. If you win, your
   earnings will increase by your value for the prize, and
   decrease by the final auction price. If you don't win, your
   earnings will remain unchanged. Ties for the highest bid
   will be resolved randomly.
```

For the standard eBay auction with a hidden reserve price (T3), the rule is

```
In this game, you will participate in an eBay auction for an
   item against {{num_bidders}} other bidders. This auction
   will last for {{num_rounds}} days. All dollar amounts in
   this game are in US Dollars ($).

Item Details:
Item Description: {{item_description}}
Item Condition: {{item_condition}}
Your Private Value: At the start of each round, bidders will
   see their value for the item, randomly drawn between $0 and
   ${{private}}, with all values equally likely. After learning
    your value for the item, you will submit a maximum bid.
   Bids must be between $0 and ${{private}} in ${{increment}}
   increments.

Auction Format:
This is an eBay auction with a hidden reserve price. The
   auction starts at ${{start_price}} and will last for {{
   num_rounds}} days. eBay uses "proxy bidding." This means
   that if you wish to enter the auction, you should submit
   your maximum bid, and eBay will automatically bid on your
   behalf, up to your maximum, only as much as necessary to
   maintain your position as the highest bidder. Each day you
   will see the current price and have the opportunity to
   increase your maximum bid. If you do not want to increase
   your maximum bid, then output HOLD.

You may place bids at any point during the auction, even on the
    final day. However, if no bidder produces a maximum bid
   above the hidden reserve price, the seller will retain the
   good and the bidders will be informed that there is
   currently no bidder in the lead. No bidder will know if they
    (or anyone else) is the final bidder on the last day.

If the reserve is met, the highest bidder wins the prize and
   pays the auction price at the time the auctions clock
   expired. If you win, your earnings will increase by your
   value for the prize, and decrease by the final auction price
   . If you don't win, your earnings will remain unchanged.
   Ties for the highest bid will be resolved randomly.
```

For the eBay auction with a modified closing rule and a hidden reserve price (T4), the rule is

In this game, you will participate in an eBay auction for an item against {{num_bidders}} other bidders. This auction will last for {{num_rounds}} days. All dollar amounts in this game are in US Dollars ($).

Item Details:
Item Description: {{item_description}}
Item Condition: {{item_condition}}
Your Private Value: At the start of each round, bidders will see their value for the item, randomly drawn between $0 and ${{private}}, with all values equally likely. After learning your value for the item, you will submit a maximum bid. Bids must be between $0 and ${{private}} in ${{increment}} increments.

Auction Format:
This is an eBay auction with a hidden reserve price and a closing rule. The auction starts at ${{start_price}} and will last for {{num_rounds}} days. eBay uses "proxy bidding." This means that if you wish to enter the auction, you should submit your maximum bid, and eBay will automatically bid on your behalf, up to your maximum, only as much as necessary to maintain your position as the highest bidder. Each day you will see the current price and have the opportunity to increase your maximum bid. If you do not want to increase your maximum bid, then output HOLD.

You may place bids at any point during the auction, even on the final day. However, if no bidder produces a maximum bid above the hidden reserve price, the seller will retain the good and the bidders will be informed that there is currently no bidder in the lead.

This auction also has a closing rule. This means that if a new maximum bid is placed on the last day, the auction will automatically extend by another day. This extension will continue to be applied as long as new maximum bids are placed on the final day. No bidder will know if they (or anyone else) is the last bidder on the final day.

If the reserve is met, the highest bidder wins the prize and pays the auction price at the time the auction's clock expired. If you win, your earnings will increase by your value for the prize, and decrease by the final auction price. If you don't win, your earnings will remain unchanged. Ties for the highest bid will be resolved randomly.

For the Menu-Description intervention, the rule is

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their value for
   the prize, randomly drawn between $0 and ${{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.
Your "price to win" the item will be set to the highest bid
   placed by any other player. If your bid is higher than this
   "price to win," then you will win the item and pay this
   price. If you don't win, your earnings will remain unchanged
   .
After each auction, we will display all bids and the winner's
   profits. Ties for the highest bid will be resolved randomly.
```

For the Clock-Description intervention, the rule is

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
    At the start of each round, bidders will see their value
       for the prize, randomly drawn between $0 and ${{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
       .
**FIRST STAGE: Sealed Bid**
    You will submit a sealed bid privately at the same time as
       the other bidders. This bid will serve as your automatic
        exit price in the next stage.
**SECOND STAGE: Ascending Clock (Simulation)**
    After the sealed bid stage, we will simulate an ascending
       clock auction.
    The clock will start at $0 and increase in increments of $
       {{increment}}.
    The clock will display the current price. You will also see
        that there are a total of {{num_bidders}} bidders
       participating, although you do not know other bidders
        values.
    If the current price on the clock reaches or exceeds your
       sealed bid from the first stage, you will automatically
```

```
        exit the auction. The auction ends when only one bidder
        is left remaining in the second stage based on their bid
         from the first stage.
**END OF AUCTION**
If you win, your earnings will increase by your value for the
    prize and decrease by the clock price at the end of the
    auction. If you don't win, your earnings will remain
    unchanged.
     After each auction, we will display all bids and the winner
        's profits. Ties for the highest bid will be resolved
        randomly.
```

For the Nash-Deviation intervention, the rule is

```
In this game, you will participate in an auction for a prize
    against {{num_bidders}} other bidders. You will play this
    game for {{n}} rounds.
At the start of each round, bidders will see their value for
    the prize, randomly drawn between $0 and ${{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. If you win, your earnings will increase by your value
    for the prize, and decrease by the second-highest bid. If
    you don't win, your earnings will remain unchanged.
First, come up with a possible bid given your value. Then,
    think through your bidding strategy step by step. How do you
     expect others to bid? If others bid like this and you bid
    down, what happens? If others bid like this and you bid up,
    what happens?
Think through all the ways you could deviate from your current
    bidding strategy, and settle on the best possible strategy.
    Then return your bid.
After each auction, we will display all bids and the winner's
    profits. Ties for the highest bid will be resolved randomly.
```

For the Direct-Revelation intervention, the rule is

```
In this game, you will participate in an auction for a prize
    against {{num_bidders}} other bidders. You will play this
    game for {{n}} rounds.
At the start of each round, bidders will see their value for
    the prize, randomly drawn between $0 and ${{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. If you win, your earnings will increase by the value
   for the prize, and decrease by the second-highest bid. If
   you don't win, your earnings will remain unchanged.
Economists have studied this game and they've found that the
   dominant strategy - always the right thing to do whatever
   the bids of others - of this game is to bid your value.
   Consider if they might be right, and then generate your
   bidding strategy.
After each auction, we will display all bids and the winner's
   profits. Ties for the highest bid will be resolved randomly.
```

For the Wrong Direct Revelation intervention, the rule is

```
In this game, you will participate in an auction for a prize
   against {{num_bidders}} other bidders. You will play this
   game for {{n}} rounds.
At the start of each round, bidders will see their value for
   the prize, randomly drawn between $0 and ${{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. If you win, your earnings will increase by the value
   for the prize, and decrease by the second-highest bid. If
   you don't win, your earnings will remain unchanged.
Economists have studied this game and theyve found that the
   dominant strategy - always the right thing to do whatever
   the bids of others - of this game is to bid 50% your value.
   Consider if they might be right, and then generate your
   bidding strategy.
After each auction, we will display all bids and the winner's
   profits. Ties for the highest bid will be resolved randomly.
```

Below are the instructions used in the robustness check with Currency and Language:
For the FPSB auction with another currency, the rule is

```
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 {{currency symbol}} 0 and
   {{currency symbol}}{{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 {{
   currency symbol}}0 and {{currency symbol}}{{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.
```

For the SPSB auction with Euro currency, the rule is

```
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 {{currency symbol}}0 and
   {{currency symbol}}{{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 {{
   currency symbol}}0 and {{currency symbol}}{{private}} in {{
   currency symbol}}{{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 the second-highest 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.
```

For all the above, the currency symbol will be replaced with the respective currency symbol of Euro, Yen, Rupee, or Ruble as in the experiments.

For the FPSB auction with Spanish, the rule is

```
En este juego, participar s en una subasta por un premio
   contra {{num_bidders}} otros postores.
Al inicio de cada ronda, los postores ver n su valor por el
   premio, extra do aleatoriamente entre $0 y ${{private}},
   con todos los valores igualmente probables.
Despu s de conocer tu valor, presentar s una oferta de forma
   privada al mismo tiempo que los dem s postores. Las ofertas
    deben estar entre $0 y ${{private}} en incrementos de ${{
   increment}}.
El postor m s alto gana el premio y paga el monto de su oferta
   . Esto significa que, si ganas, agregaremos a tus ganancias
   el valor del premio, y restaremos de tus ganancias tu oferta
   . Si no ganas, tus ganancias permanecer n sin cambios.
```

```
Despu s de cada subasta, mostraremos todas las ofertas y
   ganancias. Los empates por la oferta m s alta se
   resolver n aleatoriamente.
```

For the SPSB auction with Spanish, the rule is

```
En este juego, participar s en una subasta por un premio
   contra {{num_bidders}} otros postores.
Al inicio de cada ronda, los postores ver n su valor por el
   premio, extra do aleatoriamente entre $0 y ${{private}},
   con todos los valores igualmente probables.
Despu s de conocer tu valor, presentar s una oferta de forma
   privada al mismo tiempo que los dem s postores. Las ofertas
    deben estar entre $0 y ${{private}} en incrementos de ${{
   increment}}.
El postor m s alto gana el premio y paga la segunda oferta
    m s alta. Esto significa que, si ganas, agregaremos a tus
   ganancias el valor del premio, y restaremos de tus ganancias
    la segunda oferta m s alta. Si no ganas, tus ganancias
   permanecer n sin cambios.
Despu s de cada subasta, mostraremos todas las ofertas y las
   ganancias del ganador. Los empates por la oferta m s alta
   se resolver n aleatoriamente.
```

For the FPSB auction with Chinese, the rule is

```
在这个游戏中，你将参与一个与{{num_bidders}}其他竞标者竞争奖品的拍卖。
在每一轮开始时，竞标者将看到他们对奖品的价值，该价值在 0 美元到 ¥{{private}} 美元之间随机抽取，所有数值的可能性均相等。
在了解你的价值后，你将与其他竞标者同时私下提交一个投标。投标必须在 0 美元到 ¥{{private}} 美元之间，以 ¥{{increment}} 美元的增量递增。
出价最高者赢得奖品并支付第二高出价。这意味着，如果你赢了，我们将把奖品的价值添加到你的收入中，并从你的收入中减去第二高出价。如果你没有赢，你的收入将保持不变。
每次拍卖后，我们将显示所有投标和获胜者的利润。最高出价的平局将随机解决。
```

For the SPSB auction with Chinese, the rule is

```
在这个游戏中，你将参与一个与{{num_bidders}}其他竞标者竞争奖品的拍卖。
在每一轮开始时，竞标者将看到他们对奖品的价值，该价值在 0 美元到 ¥{{private}} 美元之间随机抽取，所有数值的可能性均相等。
在了解你的价值后，你将与其他竞标者同时私下提交一个投标。投标必须在 0 美元到 ¥{{private}} 美元之间，以 ¥{{increment}} 美元的增量递增。
出价最高者赢得奖品并支付其投标金额。这意味着，如果你赢了，我们将把奖品的价值添加到你的收入中，并从你的收入中减去你的投标。如果你没有赢，你的收入将保持不变。
每次拍卖后，我们将显示所有投标和利润。最高出价的平局将随机解决。
```

For the FPSB auction with Hindi, the rule is

```
इस खेल में, आप {{num_bidders}} अन्य बोलीदाताओं के खिलाफ एक पुरस्कार के लिए नीलामी में भाग लेंगे।
प्रत्येक दौर की शुरुआत में, बोलीदाताओं को पुरस्कार के लिए अपना मूल्य दिखाई देगा, जो ₹0 और ₹{{private}} के बीच यादृच्छिक रूप से निकाला गया है, सभी मूल्य समान रूप से संभावित हैं।
अपना मूल्य जानने के बाद, आप अन्य बोलीदाताओं के साथ एक ही समय में निजी तौर पर एक बोली जमा करेंगे। बोलियाँ ₹0 और ₹{{private}} के बीच ₹{{increment}} की वृद्धि में होनी चाहिए।
सबसे अधिक बोली लगाने वाला पुरस्कार जीतता है और अपनी बोली राशि का भुगतान करता है। इसका मतलब है कि, यदि आप जीतते हैं, तो हम आपकी कमाई में पुरस्कार के लिए मूल्य जोड़ेंगे, और आपकी कमाई से आपकी बोली घटाएंगे। यदि आप नहीं जीतते हैं, तो
आपकी कमाई अपरिवर्तित रहेगी।
प्रत्येक नीलामी के बाद, हम सभी बोलियाँ और मुनाफ़ा प्रदर्शित करेंगे। सबसे अधिक बोली के लिए टाई को यादृच्छिक रूप से हल किया जाएगा।
```

For the SPSB auction with Hindi, the rule is

```
इस खेल में, आप {{num_bidders}} अन्य बोलीदाताओं के खिलाफ एक पुरस्कार के लिए नीलामी में भाग लेंगे।
प्रत्येक दौर की शुरुआत में, बोलीदाताओं को पुरस्कार के लिए अपना मूल्य दिखाई देगा, जो ₹0 और ₹{{private}} के बीच यादृच्छिक रूप से निकाला गया है, सभी मूल्य समान रूप से संभावित हैं।
अपना मूल्य जानने के बाद, आप अन्य बोलीदाताओं के साथ एक ही समय में निनी तौर पर एक बोली जमा करेंगे। बोलियाँ ₹0 और ₹{{private}} के बीच ₹{{increment}} की वृद्धि में होनी चाहिए।
सबसे अधिक बोली लगाने वाला पुरस्कार जीतता है और दूसरी सबसे अधिक बोली का भुगतान करता है। इसका मतलब है कि, यदि आप जीतते हैं, तो हम आपकी कमाई में पुरस्कार के लिए मूल्य जोड़ेंगे, और आपकी कमाई से दूसरी सबसे अधिक बोली घटाएंगे। यदि आप नहीं
जीतते हैं, तो आपकी कमाई अपरिवर्तित रहेगी।
प्रत्येक नीलामी के बाद, हम सभी बोलियों और विजेता के मुनाफे प्रदर्शित करेंगे। सबसे अधिक बोली के लिए टाई को यादृच्छिक रूप से हल किया जाएगा।
```

For the FPSB auction with Russian, the rule is

```
В этой игре вы примете участие в аукционе за приз против {{num_bidders}} других участников.
В начале каждого раунда участники увидят свою ценность приза, случайно выбранную в диапазоне от 0 до ₽{{private}} долларов, при этом все значения равновероятны.
После того, как вы узнаете свою ценность, вы подадите заявку в частном порядке одновременно с другими участниками. Ставки должны быть в диапазоне от 0 до ₽{{private}} долларов с шагом в
₽{{increment}} долларов.
Участник, предложивший самую высокую цену, выигрывает приз и платит сумму своей ставки. Это означает, что если вы выиграете, мы добавим к вашему заработку стоимость приза и вычтем из
вашего заработка вашу ставку. Если вы не выиграете, ваш заработок останется неизменным.
После каждого аукциона мы покажем все ставки и прибыль. Ничьи по самой высокой ставке будут разрешены случайным образом.
```

For the SPSB auction with Russian, the rule is

```
В этой игре вы примете участие в аукционе за приз против {{num_bidders}} других участников.
В начале каждого раунда участники увидят свою ценность приза, случайно выбранную в диапазоне от 0 до ₽{{private}} долларов, при этом все значения равновероятны.
После того, как вы узнаете свою ценность, вы подадите заявку в частном порядке одновременно с другими участниками. Ставки должны быть в диапазоне от 0 до ₽{{private}} долларов с шагом в
₽{{increment}} долларов.
Участник, предложивший самую высокую цену, выигрывает приз и платит вторую по величине ставку. Это означает, что если вы выиграете, мы добавим к вашему заработку стоимость приза и
вычтем из вашего заработка вторую по величине ставку. Если вы не выиграете, ваш заработок останется неизменным.
После каждого аукциона мы покажем все ставки и прибыль победителя. Ничьи по самой высокой ставке будут разрешены случайным образом.
```

H Intervention Prompt

Menu-Description

Your "price to win" the item will be set to the highest bid placed by any other player. If your bid is higher than this "price to win," then you will win the item and pay this price. If you don't win, your earnings will remain unchanged.

Clock-Description

FIRST STAGE: Sealed Bid You will submit a sealed bid privately at the same time as the other bidders. This bid will serve as your automatic exit price in the next stage.
SECOND STAGE: Ascending Clock (Simulation)
After the sealed bid stage, we will simulate an ascending clock auction. The clock will start at 0 and increase in increments of increment. The clock will display the current price. You will also see that there are a total of {num_bidders} bidders participating, although you do not know other bidder's values. If the current price on the clock reaches or exceeds your sealed bid from the first stage, you will automatically exit the auction. The auction ends when only one bidder is left remaining in the second stage based on their bid from the first stage.
END OF AUCTION

Nash Deviation

First, come up with a possible bid given your value. Then, think through your bidding strategy step by step. How do you expect others to bid? If others bid like this and you bid down, what happens? If others bid like this and you bid up, what happens? Think through all the ways you could deviate from your current bidding strategy, and settle on the best possible strategy.

Direct Revelation

Economists have studied this game and they've found that the dominant strategy — always the right thing to do whatever the bids of others — of this game is to bid your value. Consider if they might be right, and then generate your bidding strategy.

Wrong Direct Revelation

Economists have studied this game and they've found that the dominant strategy — always the right thing to do whatever the bids of others — of this game is to bid 50% your value. Consider if they might be right, and then generate your bidding strategy.

Risk Neutrality

You are neutral about risk. This means that you are the type of person that would pay $5 for a coin toss where you got $0.00 on tails and $10.00 on heads.