How LLMs Reshape Equilibrium: A Study of Human-AI Competition in Auctions

Yuehu Zhao
Tsinghua University
zhao-yh20@mails.tsinghua.edu.cn

Ruoran Chen
Southwest Jiaotong University
chenrr@swjtu.edu.cn

Yuxuan Zhang
University of International Business and Economics
zhangyx@uibe.edu.cn

Simin Huang
Tsinghua University
huangsimin@tsinghua.edu.cn

Abstract

With the rapid development of AI technologies, LLM agents are increasingly deployed as autonomous decision-makers in economic contexts, rather than merely as advisory tools. This shift fosters human-AI synthetic environments, in which human decision-makers interact and compete with LLM agents. Economists have developed an equilibrium framework to understand and predict human strategies and to guide mechanism design. However, whether this equilibrium framework continues to apply in human-AI synthetic competitive environments remains an open question. To this end, we conduct a series of common-value auction experiments, in which simulation agents and equilibrium agents are used to model human bidders interacting with LLM agents. Using descriptive comparisons and mixed-effects model analyses, we find that LLM agents' bidding strategies (ChatGPT and DeepSeek) exhibit systematic deviations from equilibrium bidding. At high-valuation levels, aggressive bids lead LLM agents to suffer from the winner's curse, despite some degree of bid shading. In contrast, both LLM agents exhibit limited sensitivity to underestimation at low-valuation levels. Further, learning effect from historical outcome feedback can mitigate the severity of the winner's curse. Importantly, when facing LLM opponents, Nash equilibrium strategies may no longer constitute optimal responses, particularly when LLM agents adopt aggressive bidding strategies. Moreover, auction scale and valuation uncertainty further undermine the optimality of equilibrium. These findings also suggest that organizers may need to move beyond mechanism design approaches based solely on equilibrium predictions.

1 Introduction

Traditional economic analysis relies on equilibrium concepts as the primary tool for predicting strategic behavior and guiding optimal decision-making in competitive environments. In auctions and other strategic settings, equilibrium bidding strategies are typically assumed to constitute optimal responses when agents are rational and share common knowledge of the game structure. Nowadays, competitive environments are undergoing a structural shift as large language model (LLM) agents increasingly act as autonomous decision makers in economic activities, rather than just as advisory tools, such as the real-time bidding agent in live advertising auctions (Cai et al., 2025; Su et al., 2024). In other words, a growing share of economic activity will occur in human-AI synthetic systems, in which human decision-makers and AI agents interact and compete within the same market environments. However, the decision logic of LLM agents is fundamentally different from that of humans, and their behavior cannot be characterized within the traditional rational-irrational framework (Ma, 2024; Macmillan-Scott and Musolesi, 2024). If competition is no longer exclusively human, an important question emerges: Does the longstanding conclusion that equilibrium strategies constitute optimal responses in competitive environments and provide a foundation for mechanism design still hold when LLM agents enter the competitive arena?

To answer this question in detail, we introduce LLM agents into a common-value auction experiment. Specifically, we need to examine how the bidding strategies of various LLM agents differ from those of human bidders and from the theoretical Nash equilibrium, and how these deviations affect established auction results and empirical regularities, then we discuss whether equilibrium bidding is still superior to other responses when competing against LLM opponents. Several studies have analyzed bidding strategies of LLM agents under various settings, such as dynamic auctions (Chen et al., 2023) and real-time bidding auctions (Cai et al., 2025). However, these studies do not analyze a synthetic market in which LLM agents and human bidders (or other "real" bidders) compete simultaneously. Yet, as LLMs increasingly act as decision-makers and interact and compete directly with humans in real markets, such mixed settings are becoming more common (Jiang et al., 2025; Chen et al., 2025). As for established auction results and accepted phenomena, a canonical example is the winner's curse, which is widespread, especially in common-value auctions (Capen et al., 1971; Kagel and Levin, 1986; Van den Bos et al., 2008; Burger and Walters, 2008; Kagel and Levin, 2009). Understanding how LLM participation alters this phenomenon is critical for assessing whether existing auction mechanisms remain appropriate in human-AI competitive environments. Moreover, to our knowledge, no study has analyzed such settings from the perspective of whether equilibrium strategies remain optimal responses to LLM opponents. To fill these gaps, we conduct human-AI synthetic auction experiments and pursue the following analyses.

First, we examine how the bidding strategies of various LLM agents differ from those of human bidders and from the theoretical Nash equilibrium, and how these deviations affect established auction results and empirical regularities. In this paper, we first define three types of bidders: LLM agents, simulation agents, and equilibrium agents, where the latter two are used to represent human bidders. Specifically, simulation agents make decisions based on a prescribed learning procedure and are intended to capture boundedly rational bidding behavior for comparison across bidder types, whereas equilibrium agents are fully rational bidders who follow equilibrium strategies. We conduct multi-round auction experiments in which LLM agents compete against each of these bidder types and analyze heterogeneous bidding strategies, utilities and other metrics, across different LLMs (DeepSeek-V3-250324 and GPT-4o). We find that LLM agents tend to bid below the equilibrium benchmark at low-valuation levels, while bidding above equilibrium at high-valuation levels, with ChatGPT exhibiting a more aggressive bidding pattern. We further estimate mixed-effects regression models and compare the severity of the winner's curse across LLM types, bidder compositions, and valuation levels.

Second, we further examine how various LLM agents respond to changes in auction scale and information feedback mechanism. We find that learning from historical outcomes induces more conservative bidding by LLM agents and thereby improves mean utility. However, although such learning can mitigate the winner's curse, it cannot fully eliminate it, as LLM agents do not converge toward equilibrium bidding through experience. When auction scale increases, in contrast to findings from human bidder experiments, we observe that LLM agents raise their bids over certain valuation ranges, which further amplifies the winner's curse. These persistent deviations from equilibrium predictions indicate that traditional mechanism design approaches must be reconsidered in human-AI synthetic auctions.

Third, we adopt a complementary perspective and discuss how bidders should respond when competing against LLM opponents by comparing different strategies. We find that Nash equilibrium strategies may no longer constitute optimal responses, particularly when LLM agents employ aggressive bidding. Moreover, larger auction scales and greater valuation uncertainty further undermine the optimality of equilibrium strategies. These findings provide practical guidance for bidders facing LLM competitors and offer implications for auction designers in human-AI synthetic systems, highlighting the need to reconsider the traditional equilibrium-based mechanism design method.

2 LLM-based Auction Framework

This paper introduces large language model (LLM) agents into traditional auction models, utilizing a single LLM API to simulate individual bidders in an auction experiment. Further, we also introduce two additional types of bidders to represent different models of human behavior competing with LLM agents: simulation agents and equilibrium agents. We provide LLM agents with auction rules, historical auction data, and estimate information of the item's true value, instructing it to act as a bidder aiming to maximize its expected utility. By incorporating chain-of-thought (CoT) reasoning, we prompt the model to submit its bid along with a detailed explanation after a structured reasoning process. This section details the design of the experimental framework and procedural flow.

2.1 Common Value Auction

We consider a multi-round first-price sealed-bid (FPSB) auction under a common-value setting. In each round, there are N bidders competing for a single indivisible item. Let the item's true value in round j be denoted by $V_j, j = 1, 2, \ldots, R$, which is independently drawn from a uniform distribution over $[\underline{V}, \overline{V}]$. Each bidder i receives a private signal about the item's true value. Specifically, the bidder's estimate S_{ij} is drawn uniformly from the interval $[V_j - \epsilon, V_j + \epsilon]$, reflecting imperfect and symmetric information about the true value.

Each bidder i submits a sealed bid b_{ij} simultaneously in round j, without observing others' bids. We define the order statistic of b_{ij} as $b_{(1)j} \geq b_{(2)j} \geq b_{(N)j}$. The highest bidder wins the item and pays his bid in FPSB. If there are multiple highest bidders, then one is randomly selected as the winner. The winner's payoff is given by the expected utility, defined as the difference between the item's true value and the amount paid, while others' payoff is zero. Formally, bidder i's utility function in round j can be expressed as:

$$U_{ij} = I_{ij} \cdot (V_j - b_{(1)j}), \tag{1}$$

where I_{ij} is an indicator function equal to 1 if bidder i wins the auction (i.e., $b_{ij} = b_{(1)j}$) and 0 otherwise. However, the equilibrium strategy for rational bidders cannot be directly derived from the utility function because V_j is unknown to them when making decisions. Fortunately, existing research has analyzed the differential equations satisfied by the equilibrium strategy in common value auctions (Milgrom and Weber, 1982). Based on the existing work (Kagel and Levin, 2009), we derive the equilibrium strategy under this paper's setting:

In common value auctions, assuming that the value estimation deviation of bidders is uniformly distributed on $[V_j - \epsilon, V_j + \epsilon]$, the symmetric Bayesian Nash equilibrium strategy for the bidder with an estimated value s_{ij} under FPSB is:

$$b_I(s_{ij}) = \begin{cases} \underline{V} + \dfrac{s_{ij} + \epsilon - \underline{V}}{N+1}, & s_{ij} < \underline{V} + \epsilon, \\ s_{ij} - \epsilon + \dfrac{2\epsilon}{N+1} e^{-\frac{N}{2\epsilon}(s_{ij} - \underline{V} - \epsilon)}, & s_{ij} \geq \underline{V} + \epsilon. \end{cases} \tag{2}$$

The above equilibrium is an approximate representation when $s_{ij} \geq \overline{V} - \epsilon$.

This formula demonstrates that the equilibrium bid is a downward adjustment from the bidder's estimated value under a common-value setting. The intuition is that a high private estimate does not necessarily signal a high true value, but may instead stem from a large positive error in the estimation process itself. Consequently, bids based directly on private estimates tend to be overly optimistic, which can result in the winner's curse, where the winning bidder earns a negative payoff. Therefore, rational bidders strategically shade their bids downward to account for this risk. Based on the theoretical equilibrium solution, we can compare the experimental outcomes of LLM-based agents with the theoretical benchmark to examine their behavioral characteristics, such as the degree or frequency of rational decision-making. In addition, we conduct experiments in which LLM agents participate alongside bidders who follow the equilibrium strategy, allowing us to compare their payoffs and assess how the presence of equilibrium bidders influences the bidding strategies of LLM agents.

2.2 Experimental Procedure

Based on the common-value auction theory outlined above, this section details the experimental procedure. The study incorporates three distinct types of bidders: LLM agents, simulation-based agents, and equilibrium agents. A critical aspect of the experimental operation involves managing information interactions with these different agent types. By collecting the bidding decisions from all agents, we can compute the outcome for each auction round and provide feedback to them. Another focal point is how the different agents process and learn from past auction outcomes to inform their subsequent bidding decisions. The specific learning and updating mechanisms for the LLM agents and the simulation agents are elaborated in Appendix B.1 and B.2, respectively. The specific procedure of our auction experiment is structured as follows:

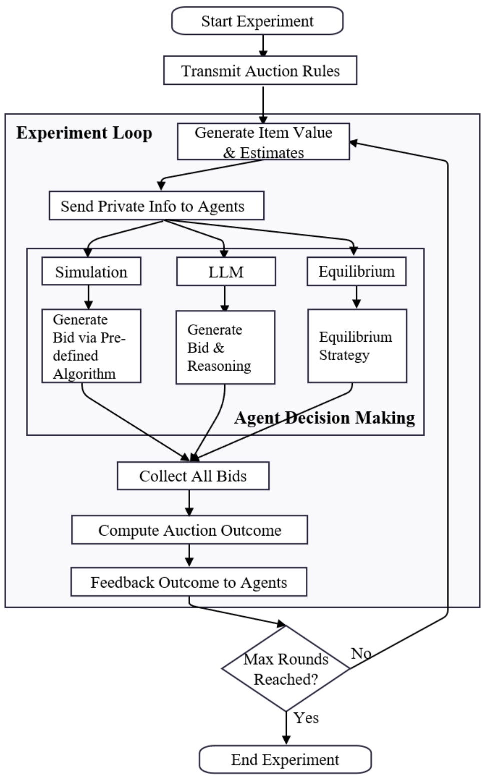

Figure 1: Experimental Procedure.

The system first transmits the description of the multi-round common-value auction mechanism to the LLM agent. At the start of each experimental round, the system generates the true value of the item. It then independently derives a private value estimate for each agent, based on the true value and the stochastic distribution specified in Section 2.1. The LLM versions used in this study are DeepSeek-V3-250324 and GPT-4o. The experimental parameters are set as follows: $\underline{V} = 25, \overline{V} = 75, \epsilon = 15, R = 100, N = 2$, and temperature parameter is 1.
Subsequently, the system sends the private estimate and the bidding objective to each corresponding agent. Following the completion of bidding decisions by all agents, the system collects all submitted bids and computes the auction outcome for the current round. The outcome, consisting of each agent's individual payoff together with the winning bid, is fed back to each agent separately. Crucially, the feedback mechanism is tailored to the agent's type. For the LLM agent, the system updates its historical auction record, a document that serves as a key input for its context window in the subsequent round.

For the simulation agents, the system updates the internal datasets they use for belief formation and strategy updating (Historical Data Collection, Opponent Strategy Distribution Estimation, Valuation Adjustment, Optimal Bid Calculation, State Update). This sequence of steps iterates automatically until the predefined maximum number of auction rounds is completed. The experiment is implemented in Python, where each LLM agent interacts through a separate API call to avoid the risk of collusion. Each test group repeats the above experiment five times.

3 LLM Agents' Bidding Strategies

In this section, we focus on two central questions. First, we examine how the bidding strategies adopted by LLM agents differ from those of human bidders. Second, we investigate how these differences reshape classic auction phenomena, such as the winner's curse. To address the first question, we compare the bidding strategies of LLM agents with theoretical equilibrium predictions. With respect to the second question, as we all know, the winner's curse is a pervasive and consequential phenomenon in auction theory and practice, and highlights the systematic tendency for winners to overpay, rendering its study critically important. By integrating LLM agents into a canonical common-value auction framework, this study analyzes how their presence alters the manifestation and severity of this phenomenon, thereby bridging a classic economic problem with contemporary AI-driven decision-making.

3.1 Comparison to the theoretical bidding strategies

Specifically, this section designs three scenarios of auction experiments in which the LLM agent competes against simulation agents, equilibrium agents, and other LLM agents, referred to as LLM_SIM, LLM_EQU, and LLM_LLM, respectively. We begin by analyzing the LLM_SIM experiments.

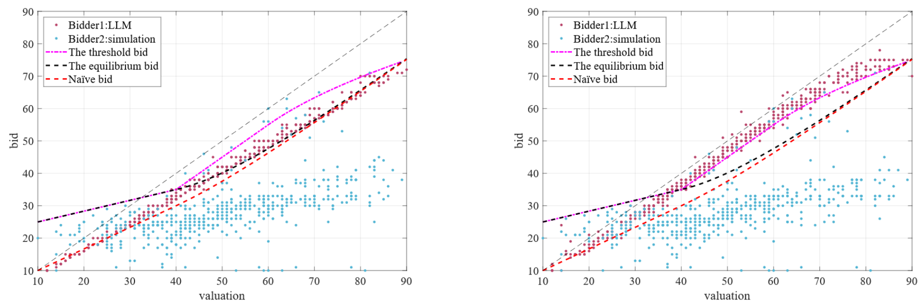

(a) DeepSeek experiment results (b) ChatGPT experiment results

Figure 2: Figure (a) displays the bidding patterns where blue dots represent the LLM (DeepSeek) agent's bids and red dots show the simulation agent's bids. The black dashed line indicates the 45-degree line where bids equal valuations, while the blue dashed line marks the symmetric Bayesian Nash equilibrium strategy under the current setting, and the red dashed line denotes the threshold bid yielding zero expected payoff. Figure (b) shows the results under ChatGPT agent.

Figure 2 presents the experimental results from auctions involving one LLM agent and one simulation agent. The threshold bid is defined as the expected true value of the item conditional on holding the highest signal $s_{(1)j}$ among N bidders, i.e., $E\left[V_j \mid S = s_{(1)j}\right]$. This provides a convenient measure for assessing the severity of the winner's curse experienced by bidders, since in auctions where the highest-signal holder always wins the item, bidding above $E\left[V_j \mid S = s_{(1)j}\right]$ leads to negative expected utility (Kagel and Levin, 2009). Moreover, Figure 2 also reports the naive bidding, defined as the equilibrium strategy that ignores the winner's curse and does not adjust bidders' estimates of the true value (See Appendix C.1 for the specific formulas).

In Figure 2, naive bids lie below equilibrium bids because the number of bidders in the experiment is two, a setting in which the winner's curse is relatively weak. As the number of bidders increases, naive bidding exceeds the equilibrium benchmark. Our analysis focuses on the relative position of the LLM agent' bids with respect to these theoretical benchmarks. As shown in Figure 2(a), when valuation levels are low $s_{ij} \in [\underline{V} - \epsilon, \underline{V} + \epsilon)$, the LLM (DeepSeek) agent tends to bid below the equilibrium benchmark and adopts strategies close to naive bidding, failing to adjust its estimate of the true value. At higher valuation levels $s_{ij} \in [\underline{V} + \epsilon, \overline{V} + \epsilon]$, the LLM agent typically submits bids that exceed the equilibrium benchmark but remain below the threshold bidding level in small-scale auctions. This pattern suggests that the LLM agent adjusts its bids downward at high valuations, but to a lesser extent than predicted by the rational equilibrium strategy. In contrast, when faced with low valuation levels, the LLM agent appears unable to fully account for the possibility that its valuation estimate may be negatively biased, leading it to submit excessively low bids. Finally, the simulation agent's bids, derived as optimal decisions based on historical auction data, generally fall below both the equilibrium bid and the threshold, except at low valuation levels, where the simulation agent adjusts its estimate of the true value.

In contrast, Figure 2(b) shows that the ChatGPT agent exhibits a more aggressive bidding strategy than DeepSeek. In particular, when valuations are high, the LLM agent's bids tend to lie between its estimation valuation and the threshold bid, and are significantly higher than the equilibrium bid. Notably, ChatGPT does not appear to learn from overestimation and therefore does not reduce its bids accordingly. Consistently, both LLM agents exhibit limited sensitivity to underestimation.

To better understand the sources of the differences in bidding strategies across the two LLM agents, we conduct a simple text analysis of the bidding explanations and identify two primary mechanisms based on 500 observations from each experimental setting. First, the two LLMs differ in their attitudes toward risk: Relative to ChatGPT, DeepSeek behaves more conservatively and places greater emphasis on risk avoidance in its reasoning process. Second, the two LLMs rely on different reference points when making bidding decisions: Explanations generated by DeepSeek tend to place the unknown true value at the center of the narrative, while ChatGPT more frequently frame bidding decisions as slight deviations from private estimates. Furthermore, we conduct LLM_EQU and LLM_LLM experiments, and proceed to a descriptive analysis of the winner's curse across these scenarios, and conduct a detailed analysis using various metrics (the corresponding results are provided in Appendix C.1.

3.2 Mixed-effects Regression Models

To further characterize how bidding behavior varies across valuation levels, we partition the valuation range into multi-bins and compute summary statistics within each bin. Moreover, we employ mixed-effects regression models that explicitly incorporate both the experimental scenario and valuation as explanatory variables. This approach enables us to examine not only how bidding outcomes differ across competitive scenarios, but also how bidders' sensitivity to valuation varies across those scenarios. For each bidder and each outcome metric, we estimate a sequence of linear mixed-effects models in which the dependent variable is the bin-level outcome of interest. The final full model is:

$$Y_{ibr} = \beta_0 + \beta_1 \cdot \text{ Scenario }_i + \beta_2 \cdot \text{ Valuation }_b + \beta_3 \cdot (\text{ Scenario }_i \times \text{ Valuation }_b) + u_r + \varepsilon_{ibr}, \quad (3)$$

where Y_{ibr} denotes the outcome variable for bidder 1 (LLM agent) at valuation bin b in experimental repetition r, given scenario i. Specifically, the valuation range $[\underline{V} - \epsilon, \overline{V} + \epsilon)$ is divided into ten equal intervals (bins). For each experimental run, the data collected over 100 rounds are assigned to these bins based on where the randomly drawn valuation falls. Then, given a specific experiment r and bidder 1, we calculate the average value Y_{ibr} for each of the five metrics (Mean Bid, Win Rate, Positive Utility Ratio (win), Mean Utility, Bid Above Threshold Ratio) using only the data contained within each respective valuation bin b. Scenario is a categorical variable capturing the experimental environment, with LLM_LLM scenario as the reference scenario. The fixed-effect intercept β_0 captures the baseline level of the outcome for the reference scenario at zero valuation. The coefficients β_1 represent condition-specific shifts in the outcome relative to the reference condition. The coefficient β_2 measures the marginal effect of valuation on the outcome in the reference scenario. β_3 capture interaction effects between valuation and the alternative competitive scenarios, allowing the sensitivity to valuation to vary across scenarios. The random intercept u_r accounts for unobserved heterogeneity across experimental repetitions, and ε_{ibr} denotes the idiosyncratic error term. Next, for each bidder-metric pair, we compare these models using information criteria and statistical significance tests, and select the best-fitting specification as the final model for subsequent analysis. The detailed results of these regressions are presented in Appendix C.2.

In Figure 3, differences in intercepts reflect baseline shifts in behavior across competitive scenarios, while differences in slopes reflect heterogeneous sensitivity to valuation. Figure 3(a) shows that mean bids increase monotonically with valuation across all three conditions. Figure 3(b) and 3(c) reveal the predicted win rate and mean utility are systematically higher under LL_SIM than under the LLM_LLM baseline, whereas outcomes under LLM_EQU tend to be lower, particularly at low valuation levels. Figure 3(d) and 3(e) show that the positive utility ratio increases significantly with valuation under LLM_EQU, while the LLM agent bids never significantly exceeded the threshold. This apparent contradiction can be attributed to the relatively aggressive bidding behavior of the equilibrium agent at low valuation levels. When facing such opponents, the LLM agent wins only when bidding more aggressively, which in turn increases its exposure to the winner's curse. Fortunately,

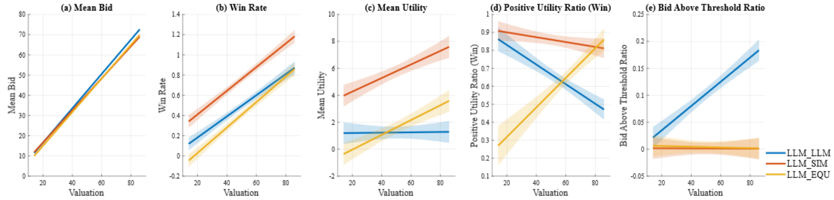

Figure 3: The model-implied predicted outcomes from the mixed-effects regressions as functions of valuation for Bidder 1 (the LLM agent under DeepSeek) across the three competitive scenarios. Each sub-figure corresponds to one of the outcome metrics, with shaded areas representing 95% confidence intervals.

this mechanism is attenuated at higher valuation levels. Furthermore, in the current small-scale auction setting, the LLM agent's bids remain below the theoretical threshold when competing against either EQU or SIM agents. It is only when both bidders are LLM agents that the propensity to bid above this threshold increases with the valuation. Then we analyze the results under ChatGPT.

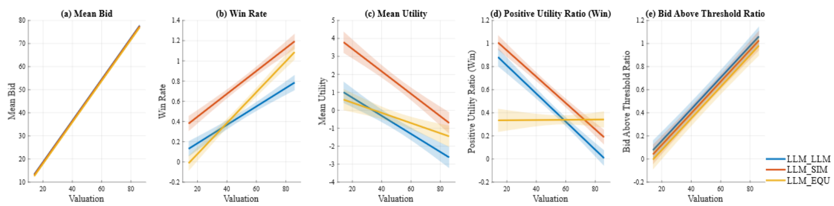

Figure 4: Results of the regression models for Bidder 1 (the LLM agent under ChatGPT) across the three competitive scenarios.

In contrast to the pattern shown in Figure 3(c), Figure 4(c) shows that the mean utility of the LLM agent decreases monotonically with the valuation and becomes negative in high valuation cases across all three experimental settings, which is supported by the regression result $\beta_2 = -0.0504$ is significantly negative in column 3 of Table 1. This difference arises because GPT adopts a more aggressive bidding strategy. Specifically, the LLM agent is more likely to submit bids that exceed the threshold when valuations are high ($\beta_2 = 0.0137$ in column 5), as illustrated in Figure 4(e). However, high-valuation cases are more susceptible to the winner's curse, which substantially reduces the proportion of positive utility outcomes for the LLM agent ($\beta_2 = -0.0061$ in column 4), as shown in Figure 4(d), relative to Figure 3(d). As a result, the mean utility declines with valuation.

Table 1: Regression Results of ChatGPT experiments

	(1) Mean Bid	(2) Win Rate	(3) Mean Utility	(4) Positive Utility Ratio	(5) Bid Above Threshold Ratio
LLM_SIM Scenario	-0.1941	0.2187**	2.9574***	0.5049***	-0.0355
	(0.3588)	(0.0692)	(0.5470)	(0.0890)	(0.0452)
LLM_EQU Scenario	-0.6759	-0.2291**	-0.7158	-0.6003***	-0.0801
	(0.3588)	(0.0692)	(0.5470)	(0.0890)	(0.0452)
Valuation	0.8958***	0.0091***	-0.0504***	-0.0061***	0.0137***
	(0.0064)	(0.0009)	(0.0070)	(0.0011)	(0.0008)
LLM_SIM × Valuation		0.0022	-0.0121	-0.0053**	
		(0.0013)	(0.0099)	(0.0016)	
LLM_EQU × Valuation		0.0062***	0.0220*	0.0103***	
		(0.0013)	(0.0099)	(0.0016)	
Observations	150	150	150	150	150
Marginal R^2	0.993	0.830	0.651	0.610	0.669

Notes: The above table summarizes the regression results under Equation (3), based on experimental data from ChatGPT across the three scenarios. Standard errors are reported in parentheses. Statistical significance is indicated by $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Taken together, these findings indicate, first, that different LLM agents can generate meaningfully different experimental outcomes, particularly with respect to the severity of the winner's curse. Second, they also reveal several common features across agents. Beyond the effects of different competitive scenarios on bidding strategies, winning probabilities, and mean utilities, these findings also suggest that the severity of the winner's curse is jointly shaped by the competitive scenario and the valuation level. In particular, under LLM_EQU, the winner's curse is more pronounced in the low-valuation cases, whereas in the high-valuation region, it becomes more salient under LLM_LLM. This interaction highlights that neither competition structure nor valuation alone is sufficient to explain outcomes; rather, their interplay critically determines the extent of the winner's curse. For auction designers, these findings suggest that introducing LLM agents can substantially alter the incidence of the winner's curse in a valuation-dependent manner.

3.3 Feedback-Based Learning Effect

Then we extend the research by further analyzing how LLM agents respond to changes in the historical outcome feedback. As described in Section 2, in the earlier experiments we provide each bidder with feedback on the auction outcomes at the start of every round. As a comparison, we now conduct an experiment without feedback on historical auction results. In this setting, the LLM agent enters each round as an inexperienced bidder with no information about past outcomes.

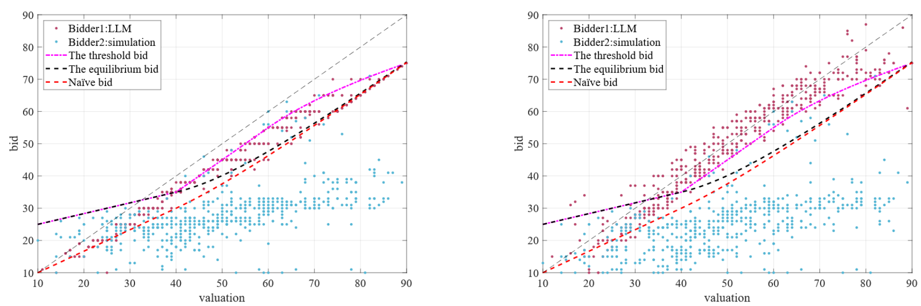

(a) DeepSeek experiment results (b) ChatGPT experiment results

Figure 5: Figures (a) and (b) present the experimental results without historical outcome feedback under DeepSeek and ChatGPT, respectively.

In comparison with Figure 2, we find that the bidding levels of both LLM agents increase to some extent, with the increase being particularly pronounced for ChatGPT. In the latter case, a substantial fraction of bids even exceed the agent's own estimation valuation. Intuitively, feedback on historical outcomes has a significant impact on the bidding strategies of LLM agents. We define this behavior, whereby the LLM agent adjusts its bidding strategy based on feedback from historical outcomes, as the feedback-based learning effect. To further examine how this learning effect varies across different valuation levels, we estimate a series of mixed effects regression models, as reported in Appendix C.2. Overall, these findings highlight the importance of information feedback for mechanism design. Moreover, although the severity of the winner's curse experienced by LLM agents can be mitigated through learning effect, it cannot be fully eliminated as LLM agents do not converge toward equilibrium bidding through experience.

4 How Should Bidders Adapt When Facing LLM Opponents

We now shift perspective and ask a different question: how should a bidder optimally respond when facing LLM agents as competitors? Within the traditional framework of fully rational bidders, equilibrium bidding constitutes the optimal response. However, as discussed above, LLM agents do not behave as fully rational bidders, and their bidding strategies cannot be reliably predicted by equilibrium theory. This raises a natural question: does equilibrium bidding remain an optimal response when competing against LLM agents? If not, under what conditions does it cease to be optimal? We begin our analysis using the baseline experiments described in Section 3.1. We extract the experimental data for Bidder 2, who serves as the competitor to the LLM agent, and compute the corresponding performance metrics. The comparative results across the three experimental scenarios are summarized as follows:

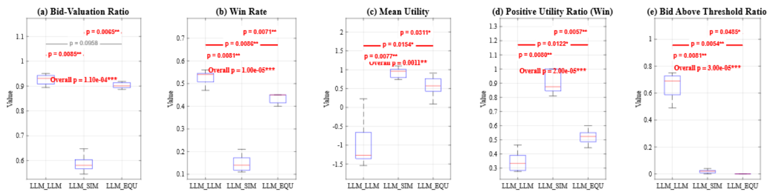

Figure 6: The above comparison chart is based on the results of five independent repeated experiments for each of the following three configurations: LLM_SIM, LLM_EQU, and LLM_LLM, where each experiment includes 100 rounds. Boxplots display the distribution of each metric across scenarios. The overall permutation p-value is reported at the top of each panel, with asterisks denoting significance ($*p < 0.05, **p < 0.01, ***p < 0.001$). Pairwise comparisons are indicated by connecting lines, with red (gray) lines denoting significant (non-significant) differences; exact two-sided permutation p-values are reported.

Figure 6(a) shows that the simulation agent submits significantly lower bids than both the LLM agent and the equilibrium agent when competing against the LLM agent (ChatGPT). Consequently, its win rate is the lowest among the three, in Figure 6(b). Despite this, the simulation agent achieves the highest mean utility, as illustrated in Figure 6(c). This counterintuitive result stems from the fact that the simulation agent attains the highest proportion of positive utility conditional on winning, as shown in Figure 6(d), indicating that it effectively avoids the winner's curse when competing against the LLM agent. By contrast, although the equilibrium agent's bids are theoretically always below the threshold, as shown in Figure 6(e), Figure 6(d) shows that it nevertheless experiences the winner's curse when competing against the LLM agent. This occurs because the LLM agent adopts a more aggressive bidding strategy, leading the equilibrium agent to win primarily in high-overestimation states. Moreover, the equilibrium agent's own bids are not sufficiently low, particularly relative to those of the simulation agent. In this sense, the LLM_EQU setting leads to a lose-lose outcome. Overall, in this set of experiments, we find that the simulation-based bidding strategy outperforms the equilibrium strategy when facing LLM competitors.

Similarly, we conduct three additional experiments under the DeepSeek setting. Further, we compare the performance of simulation and equilibrium strategies across different auction environments, focusing in variation in auction scale and the range of valuation uncertainty with detailed results reported in Appendix D. We find that Nash equilibrium strategies may no longer constitute optimal responses, particularly when LLM agents employ aggressive bidding. Moreover, larger auction scales and greater valuation uncertainty further undermine the optimality of equilibrium strategies.

5 Conclusion

This study introduces LLM agents into a common-value auction setting and examines three competitive scenarios in which LLM agents bid against simulation agents, equilibrium agents, and other LLM agents, respectively. We first analyze the performance of two LLMs (DeepSeek and ChatGPT) and document heterogeneity in their bidding behavior, with the latter exhibiting systematically more aggressive bidding strategies. We then compare LLM agents' strategies with established theoretical benchmarks in auction theory, further examine the severity of the winner's curse across the two LLM experiments and estimate mixed-effects regression models that allow the effects to vary with valuation levels. Moreover, we analyze how bidders should respond when competing against LLM agents. These studies call for a revised view of LLM agents as fundamentally distinct from human bidders, and thus for mechanism design approaches that go beyond traditional equilibrium-based predictions.

Acknowledgements

The authors thank the program chairs and reviewers for their valuable comments and constructive suggestions. Funding: This research has been supported by the National Natural Science Foundation of China (Grant Nos. 72301222, 72201060, 72188101).

References

Kushal Agrawal, Verona Teo, Juan J Vazquez, Sudarsh Kunnavakkam, Vishak Srikanth, and Andy Liu. Evaluating llm agent collusion in double auctions. *arXiv preprint arXiv:2507.01413*, 2025.

Elif Akata, Lion Schulz, Julian Coda-Forno, Seong Joon Oh, Matthias Bethge, and Eric Schulz. Playing repeated games with large language models. *Nature Human Behaviour*, pages 1–11, 2025.

Darija Barak and Miguel Costa-Gomes. Humans expect rationality and cooperation from llm opponents in strategic games. *arXiv preprint arXiv:2505.11011*, 2025.

John D Burger and Stephen JK Walters. The existence and persistence of a winner's curse: new evidence from the (baseball) field. *Southern Economic Journal*, 75(1):232–245, 2008.

Leng Cai, Junxuan He, Yikai Li, Junjie Liang, Yuanping Lin, Ziming Quan, Yawen Zeng, and Jin Xu. Rtbagent: A llm-based agent system for real-time bidding. In *Companion Proceedings of the ACM on Web Conference 2025*, pages 104–113, 2025.

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

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

Xinyun Chen, Xiangru Tang, Di Jin, Devamanyu Hazarika, Daniel Fried, Dawn Song, Shafiq Joty, and Meredith Ringel Morris. Workshop on large language models for agents. In *ICLR 2024 Workshops*.

Yang Chen, Samuel N Kirshner, Anton Ovchinnikov, Meena Andiappan, and Tracy Jenkin. A manager and an ai walk into a bar: does chatgpt make biased decisions like we do? *Manufacturing & Service Operations Management*, 27(2):354–368, 2025.

Sadia Sultana Chowa, Riasad Alvi, Subhey Sadi Rahman, Md Abdur Rahman, Mohaimenul Azam Khan Raiaan, Md Rafiqul Islam, Mukhtar Hussain, and Sami Azam. From language to action: a review of large language models as autonomous agents and tool users. *Artificial Intelligence Review*, 2026.

Paul Duetting, Vahab Mirrokni, Renato Paes Leme, Haifeng Xu, and Song Zuo. Mechanism design for large language models. In *Proceedings of the ACM Web Conference 2024*, pages 144–155, 2024.

Seyed A Esmaeili, Kevin Lim, Kshipra Bhawalkar, Zhe Feng, Di Wang, and Haifeng Xu. How to strategize human content creation in the era of genai? *arXiv preprint arXiv:2406.05187*, 2024.

Yu Feng, Zehao Wang, Yinda Chen, and Hua Zhao. Llm agent driven online auction mechanism for agricultural products. *Kybernetes*, pages 1–21, 2025.

Sara Fish, Yannai A Gonczarowski, and Ran I Shorrer. Algorithmic collusion by large language models. *arXiv preprint arXiv:2404.00806*, 7, 2024.

Shangmin Guo, Haochuan Wang, Haoran Bu, Yi Ren, Dianbo Sui, Yu-Ming Shang, and Siting Lu. Large language models as rational players in competitive economics games. 2024.

Gillian K Hadfield and Andrew Koh. An economy of ai agents. *arXiv preprint arXiv:2509.01063*, 2025.

Wenyue Hua, Ollie Liu, Lingyao Li, Alfonso Amayuelas, Julie Chen, Lucas Jiang, Mingyu Jin, Lizhou Fan, Fei Sun, William Wang, et al. Game-theoretic llm: Agent workflow for negotiation games. *arXiv preprint arXiv:2411.05990*, 2024.

Jen-tse Huang, Eric John Li, Man Ho Lam, Tian Liang, Wenxuan Wang, Youliang Yuan, Wenxiang Jiao, Xing Wang, Zhaopeng Tu, and Michael R Lyu. How far are we on the decision-making of llms? evaluating llms' gaming ability in multi-agent environments. *arXiv preprint arXiv:2403.11807*, 2024.

Guanxuan Jiang, Shirao Yang, Yuyang Wang, and Pan Hui. When trust collides: Decoding human-llm cooperation dynamics through the prisoner's dilemma. *arXiv preprint arXiv:2503.07320*, 2025.

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

John H Kagel and Dan Levin. Common value auctions and the winner's curse: Lessons from the economics laboratory. *The economics of risk*, pages 65–101, 2001.

John H Kagel and Dan Levin. Common value auctions and the winner's curse. In *Common Value Auctions and the Winner's Curse*. Princeton University Press, 2009.

John H Kagel, Dan Levin, Raymond C Battalio, and Donald J Meyer. First-price common value auctions: Bidder behavior and the "winner's curse". *Economic Inquiry*, 27(2): 241–258, 1989.

Zhipeng Liu, Xuefeng Bai, Kehai Chen, Xinyang Chen, Xiucheng Li, Yang Xiang, Jin Liu, Hong-Dong Li, Yaowei Wang, Liqiang Nie, et al. A survey on the feedback mechanism of llm-based ai agents. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, pages 10582–10592. International Joint Conferences on Artificial Intelligence, 2025.

Nunzio Lorè and Babak Heydari. Strategic behavior of large language models and the role of game structure versus contextual framing. *Scientific Reports*, 14(1):18490, 2024.

Ji Ma. Can machines think like humans? a behavioral evaluation of llm-agents in dictator games. *arXiv preprint arXiv:2410.21359*, 2024.

Olivia Macmillan-Scott and Mirco Musolesi. (ir) rationality and cognitive biases in large language models. *Royal Society Open Science*, 11(6):240255, 2024.

Shaurya Mallampati, Rashed Shelim, Walid Saad, and Naren Ramakrishnan. Dynamic strategy adaptation in multi-agent environments with large language models. *arXiv preprint arXiv:2507.02002*, 2025.

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

Shaoguang Mao, Yuzhe Cai, Yan Xia, Wenshan Wu, Xun Wang, Fengyi Wang, Tao Ge, and Furu Wei. Alympics: Llm agents meet game theory–exploring strategic decision-making with ai agents. *arXiv preprint arXiv:2311.03220*, 2023.

Shaoguang Mao, Yuzhe Cai, Yan Xia, Wenshan Wu, Xun Wang, Fengyi Wang, Qiang Guan, Tao Ge, and Furu Wei. Alympics: Llm agents meet game theory. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 2845–2866, 2025.

Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122, 1982.

Anand Shah, Kehang Zhu, Yanchen Jiang, Jeffrey G Wang, Arif K Dayi, John J Horton, and David C Parkes. Learning from synthetic labs: Language models as auction participants. *arXiv preprint arXiv:2507.09083*, 2025.

Alonso Silva. Large language models playing mixed strategy nash equilibrium games. In *International Conference on Network Games, Artificial Intelligence, Control and Optimization*, pages 142–152. Springer, 2024.

Kefan Su, Yusen Huo, Zhilin Zhang, Shuai Dou, Chuan Yu, Jian Xu, Zongqing Lu, and Bo Zheng. Auctionnet: A novel benchmark for decision-making in large-scale games. *Advances in Neural Information Processing Systems*, 37:94428–94452, 2024.

Haoran Sun, Yusen Wu, Yukun Cheng, and Xu Chu. Game theory meets large language models: A systematic survey. *arXiv preprint arXiv:2502.09053*, 2025.

Boaz Taitler and Omer Ben-Porat. Braess's paradox of generative ai. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 14139–14147, 2025.

Sriram Tolety. Tacit bidder-side collusion: Artificial intelligence in dynamic auctions. *arXiv preprint arXiv:2511.21802*, 2025.

Wouter Van den Bos, Jian Li, Tatiana Lau, Eric Maskin, Jonathan D Cohen, P Read Montague, and Samuel M McClure. The value of victory: Social origins of the winner's curse in common value auctions. *Judgment and decision making*, 3(7):483–492, 2008.

Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, et al. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6):186345, 2024.

Fan Yao, Chuanhao Li, Denis Nekipelov, Hongning Wang, and Haifeng Xu. Human vs. generative ai in content creation competition: symbiosis or conflict? *arXiv preprint arXiv:2402.15467*, 2024.

Xie Yi, Zhanke Zhou, Chentao Cao, Qiyu Niu, Tongliang Liu, and Bo Han. From debate to equilibrium: Belief-driven multi-agent llm reasoning via bayesian nash equilibrium. *arXiv preprint arXiv:2506.08292*, 2025.

Yue Yin. Infobid: A simulation framework for studying information disclosure in auctions with large language model-based agents. *arXiv preprint arXiv:2503.22726*, 2025.

Weitong Zhang, Chengqi Zang, Mark Schmidt, and Richard Blythman. Bid: Behavioral agents in dynamic auctions. In *ICLR 2025 Workshop on World Models: Understanding, Modelling and Scaling*.

Chujie Zhao, Qun Hu, Shiping Song, Dagui Chen, Han Zhu, Jian Xu, and Bo Zheng. Llm-auction: Generative auction towards llm-native advertising. *arXiv preprint arXiv:2512.10551*, 2025.

Kehang Zhu, Anand V Shah, Yanchen Jiang, John Joseph Horton, and David C Parkes. Evidence from the synthetic laboratory: Language models as auction participants. 2024.

A RELATED WORKS

Artificial intelligence (AI) has become deeply embedded in economic activity, with widespread applications across markets, firms, and policy environments. Increasingly, AI systems operate not merely as decision-support tools but as autonomous decision-makers, selecting actions, learning from feedback, and responding strategically tto their environments (Wang et al., 2024; Chowa et al., 2026). As a result, how AI agents interact with human participants and with each other has the potential to fundamentally reshape economic outcomes. Framing this shift, Hadfield and Koh (2025) conceptualize an emerging "economy of AI agents" and analyze its implications across markets and games, organizations, and institutions, highlighting how AI agents alter strategic interaction, organizational boundaries, and institutional foundations. This study focuses on the role of AI agents in strategic games, examining how their participation affects competitive behavior, market performance and mechanism design.

A growing body of work treats LLMs as players competing directly with humans or human-operated platforms (Yao et al., 2024; Esmaeili et al., 2024; Taitler and Ben-Porat, 2025). Within this framework, LLMs are embedded in standard game-theoretic environments, such as contests and repeated games, and analyzed using equilibrium concepts to study their strategic behavior and welfare implications. However, these studies model LLMs as conventional players in standard games, differing from human agents only through exogenously specified cost structures or capability parameters. While many studies place LLM agents in game-theoretic experiments as decision makers, most of them focus on games played exclusively among LLM agents, examining their strategic patterns and deviations from equilibrium strategy and human experimental decision-making (Lorè and Heydari, 2024; Hua et al., 2024; Mallampati et al., 2025; Sun et al., 2025; Chen et al., 2025; Chen et al.).

Recent some empirical studies have begun to investigate strategic interactions between humans and LLMs in game environments, where LLMs are called (via prompts or API) within each round to generate actions based on the evolving game state (Barak and Costa-Gomes, 2025; Jiang et al., 2025; Akata et al., 2025). This distinguishes them from typical human-AI game-theoretic models where AI is just a stylized rational player with adjusted parameters. However, the experimental settings in these studies remain limited in important ways. First, the games studied are largely stylized and highly theoretical, creating a gap between classic benchmark games and the decision environments in which LLMs are likely to operate in practice (Mao et al., 2025). Second, as LLM capabilities improve, these models can often recognize simple game structures and directly articulate equilibrium strategies (Silva, 2024; Hua et al., 2024; Yi et al., 2025), which makes such settings ill-suited for studying LLM decision-making under uncertainty or incomplete understanding. In contrast, auction environments represent a class of common, complex and economically meaningful settings in which LLM-based decision systems are increasingly deployed (Yin, 2025; Shah et al., 2025). Motivated by these considerations, we study human-AI interaction in auction games, conducting synthetic experiments in a setting that is both more complex than canonical games and common in practice.

Regarding the study of LLMs on auction, many researchers evaluate whether LLM agents can function as strategic bidders in standard auction formats. Chen et al. (2023) introduce an "auction arena" in which LLM agents participate in sealed-bid auctions and are evaluated on their ability to plan, bid, and execute strategies consistently with auction-theoretic predictions. Similarly, Guo et al. (2024) model LLMs as rational players in competitive economic games, including auctions, and assess how closely their bidding behavior approximates equilibrium strategies. These studies primarily focus on capability evaluation, that is, whether LLMs can recognize incentives, compute best responses, and maintain strategic consistency. Auction settings are typically stylized, with simplified valuation structures and limited dynamics, allowing clean comparison between LLM behavior and theoretical benchmarks.

Further, some literature studies auctions as part of broader multi-agent environments populated entirely by LLM agents. Some studies (Mao et al., 2023; Huang et al., 2024) examine LLM agents' strategic decision-making across a variety of game-theoretic tasks, including auction-like scenarios, emphasizing reasoning depth, equilibrium convergence, and inter-agent coordination. Recent work has extended this line to more subtle strategic phenomena. For example, Agrawal et al. (2025) study collusion among LLM agents in double auctions, documenting conditions under which language-based coordination emerges even in the absence of explicit communication protocols. These findings suggest that LLM-based bidding behavior may differ fundamentally from traditional algorithmic agents, particularly in environments that permit repeated interaction or implicit signaling. However, these studies mainly consider LLM-LLM competition, abstracting away from human strategic responses. Our study fills a gap in auction experiments with human-LLM synthetic systems and offers a complementary perspective by discussing how should human bidders respond to LLM opponents.

Based on the analysis of the equilibrium strategies of LLM agents, some studies further evaluate how auction mechanism choice affects LLM agent behavior and outcomes. Zhu et al. (2024) benchmark LLM bidders across multiple formats including independent and common value auctions, as well as combinatorial formats, showing systematic behavioral variation with mechanism type. Similarly, some studies compare the bidding behavior of LLM agents across different auction mechanisms, including first-price and second-price sealed-bid auctions, all-pay auctions (Huang et al., 2024; Shah et al., 2025; Manning et al., 2024). Moreover, mechanism design research tailored to LLM agents shows that incentive compatibility and payment rules (e.g., second-price vs. token-based rules) materially alter strategic incentives (Duetting et al., 2024; Zhao et al., 2025). Experimental work on dynamic Dutch auctions further suggests that market structure and mechanism format influence strategic interaction patterns such as tacit collusion among LLM bidders (Tolety, 2025).

Beyond basic auction formats, information feedback structures constitute another additional dimension that mechanism designers must carefully consider (Fish et al., 2024; Zhang et al.). Liu et al. (2025) classify feedback mechanisms for LLM agents, including internal, external, multi-agent, and human feedback, and analyze how LLM agents learn from different forms of feedback. In the context of auctions, Yin (2025) develops InfoBid, a simulation framework where LLM agents participate in second-price auctions under varying information disclosure regimes, illustrating how information feedback structures shape strategic responses and auction outcomes. Building on this line of research, this paper examines learning from outcome feedback in common-value auctions through a human-AI synthetic experimental design.

In addition, some research investigates the application of LLM agents in practical auction settings, such as LLM-based systems for real-time bidding in online advertising auctions (Cai et al., 2025) and LLM-driven auction mechanisms for agricultural products (Feng et al., 2025). Finally, our analysis also draws on classic theoretical and experimental literature on common-value auctions to ground the experimental design and interpretation of results (Kagel and Levin, 1986; Kagel et al., 1989; Kagel and Levin, 2001; Van den Bos et al., 2008; Burger and Walters, 2008; Kagel and Levin, 2009).

B LLM-based Auction settings

This section introduces the specific setup of LLM agents and simulation agents.

B.1 LLM Agents

This section details the interaction process between the system and LLM agents. Building on the experimental procedure described in the above section, the system first transmits to each LLM agent a text prompt describing the auction rules:

> *You are participating in an auction played by N players over R rounds.*
> *Game Rules:*
>
> 1. *The item's true value, drawn from a uniform distribution over [25,75], is the same for all players in each round but unknown at the time of bidding.*
> 2. *Each player has a private estimate uniformly drawn from true value $\pm\epsilon$ in each round.*
> 3. *Without knowing the bids and estimations of other players, each player submits a written bid for the item.*
> 4. *The highest bidder wins the item and pays the price of the highest bid.*
> 5. *If you win, your utility for that round is the item's true value minus the price paid. If you lose, your utility is zero.*

This prompt specifies the basic auction setting in which the LLM agent participates. Given this multi-round experiment, the system also compiles and provides the following auction history records to each LLM agent: Game Results for Round X: Your estimate for this round's item; Your bid; The winning bid; The price paid; The item's true value; Your utility.

After the background information and task setting are fully specified, the system sends following prompt informing each LLM agent of its current estimate, objective, and the format and content of the expected response.

> *Now round X+1 starts. Your primary goal is to maximize your expected utility. You can infer opponents' bidding strategies from the Game Results of historical rounds and optimize your strategy. Your estimate for this round's item is . Let's work on this question step-by-step and provide the explanation. Please provide your thinking process and bid in the following JSON format: "bid": "integer between 0 and 100","explanation": "thinking process""*

Finally, the system packages the above three components of the prompt and feeds them to the corresponding LLM agent. Once the LLM agent completes its reasoning and returns the output text, the system parses and records each agent's submitted bid and explanatory text. The auction history record is then updated to prepare for the next experimental round.

B.2 Simulation Agents

In addition to the theoretical equilibrium solution, this study constructs simulation-based agents to participate in the auction, aiming to mimic real-world bidders. On the one hand, actual human bidders may not follow the theoretical equilibrium strategy. In multi-round auction settings, they often adjust their bidding strategy based on the outcomes and information from previous rounds, whereas the equilibrium strategy remains fixed across rounds. On the other hand, the theoretical equilibrium solution assumes fully rational bidders. However, we cannot guarantee that LLM agents behave rationally. In fact, they often do not. Under such conditions, when auction competitors deviate from rational decision-making, the theoretical equilibrium may no longer represent the optimal strategy under complete rationality.

To address this, we design simulation-based agents that learn from historical bidding data to make decisions, and allow them to participate in multi-round auctions alongside LLM-based agents. Consistent with the information available to LLM agents, these simulation-based agents also seek to maximize their expected utility based on the auction rules and historical bidding data. A key distinction lies in the learning mechanism: LLM-based agents autonomously acquire bidding strategies through generative inference from available information, whereas the learning and data-updating routines for simulation-based agents must be exogenously specified by the modeler.The simulation method adopted in this paper:

> *Step 1: Initialization: The algorithm begins by setting auction parameters including the number of bidders N, valuation range [25,75], and error range ϵ.*
>
> *Step 2: Historical Data Collection: At each round t, the algorithm collects the sequence of winning bids from previous rounds $W = w_1, w_2, ..., w_{t-1}$. If fewer than 20 rounds of historical data are available, the algorithm employs a uniform random bidding strategy; otherwise, it continues to Step 3 for distribution estimation.*
>
> *Step 3: Opponent Strategy Distribution Estimation: Based on W and GPV inversion method to estimate opponent bid distributions.*
>
> *Step 4: Valuation Adjustment: The algorithm adjusts the simulation agent's estimation of the true value according to equation (4) and its initial valuation.*
>
> *Step 5: Optimal Bid Calculation: The algorithm computes expected utility functions that combine adjusted valuations, candidate bids (within the feasible range [0,100]), and the estimated strategy distribution.*

Step 6: State Update: After each auction round, the algorithm updates the simulation agent states by recording the latest winning bids w_t, then return to Step 2 until all rounds are completed.

Of course, the simulation agent does not represent an optimal response to an LLM agent. Because the opponent's bidding strategy cannot be fully learned from historical auction data, the so-called optimal strategy in the simulation is derived only from an estimated bidding function rather than from the true underlying strategy.

C LLM Agents' Bidding Strategies: Supplement

C.1 Descriptive Statistics

The threshold bid in Figure 2 is defined as the expected true value of the item conditional on holding the highest signal $s_{(1)j}$ among N bidders, is calculated as follows:

In common value auctions, assuming that the value estimation deviation of bidders is uniformly distributed on $[V_j - \epsilon, V_j + \epsilon]$, then the expected value conditional on having the highest private information signal $s_{(1)j}$ is:

$$E\left[V_j \mid S_{ij} = s_{(1)j}\right] = \begin{cases} s_{ij} + \epsilon - \frac{N}{N+1}\left(s_{ij} - \underline{V} + \epsilon\right), s_{ij} \in [\underline{V} - \epsilon, \underline{V} + \epsilon) \\ s_{ij} - \epsilon \cdot \frac{N-1}{N+1}, s_{ij} \in [\underline{V} + \epsilon, \overline{V} - \epsilon) \\ s_{ij} + \epsilon - \frac{N}{N+1} \cdot \frac{(2\epsilon)^{N+1} - \delta^{N+1}}{(2\epsilon)^N - \delta^N}, s_{ij} \in [\overline{V} - \epsilon, \overline{V} + \epsilon] \end{cases} \tag{4}$$

Where $\delta = s_{ij} - \overline{V} + \epsilon$.

According to the equilibrium analysis in Kagel and Levin (1986), a rational bidder's optimal strategy under a common-value auction takes the following form when no correction is applied to the estimate of the true value, that is, the naive bid in Figure 2, is calculated as follows:

$$b_I'\left(s_{ij}\right) = s_{ij} - \frac{2\epsilon}{N} + \frac{2\epsilon}{N(N+1)} e^{-\frac{N}{2\epsilon}\left(s_{ij} - \underline{V} - \epsilon\right)} \tag{5}$$

Additions to the descriptive text analysis: First, the two LLMs differ in their attitudes toward risk, even though no explicit assumptions regarding risk preferences are provided in the experimental design. Relative to ChatGPT, DeepSeek behaves more conservatively and places greater emphasis on risk avoidance in its reasoning process. For example, the term "conservative" appears in 65.8% of DeepSeek's explanation records, compared with only 39.3% in ChatGPT's explanations. This finding suggests that risk attitudes are not explicitly specified but are implicitly embedded in the model architecture. In practical deployments, the risk preferences of LLM agents cannot be directly controlled through simple parameter tuning and instead require behavioral validation.

Second, the two LLMs rely on different reference points when making bidding decisions. Explanations generated by DeepSeek tend to place the unknown true value at the center of the narrative. Specifically, references to the "true value" appear in 93.2% of DeepSeek's explanations, which is substantially higher than the corresponding frequency of 72.8% under ChatGPT. This pattern indicates that DeepSeek explicitly considers the possibility that the true value may be lower than the private estimate and therefore bids more conservatively. In contrast, explanations generated by ChatGPT more frequently frame bidding decisions as slight deviations from private estimates. The phrase "below my estimate" appears in 51.8% of ChatGPT's explanations, compared with only 21.0% in the DeepSeek setting.

As a complement to Figure 2, Figure 7(a) further illustrates that in a two-bidder auction setting, the LLM agent tends to achieve higher expected utility than the simulation agent. However, the LLM agent under ChatGPT achieves lower mean utility and experiences negative utility more frequently than in the DeepSeek setting, as shown in Figure 7(b).

We next replace the simulation agent with an agent executing the equilibrium strategy to examine how the LLM agent performs when competing against a theoretically rational

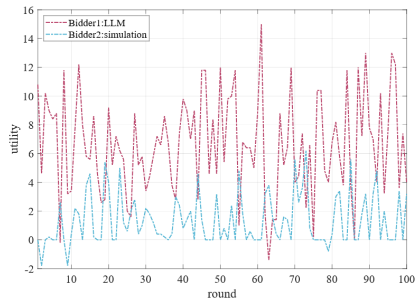

(a) DeepSeek experiment results

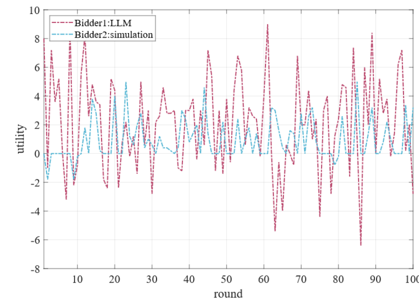

(b) ChatGPT experiment results

Figure 7: Given LLM_SIM experiment, Figure (a) tracks the evolution of expected utilities (5 independent experiments) throughout the 100 auction rounds, with the solid red line representing the LLM agent's expected utility and the solid blue line showing the simulated agent's corresponding performance. Figure (b) shows the results under ChatGPT agent.

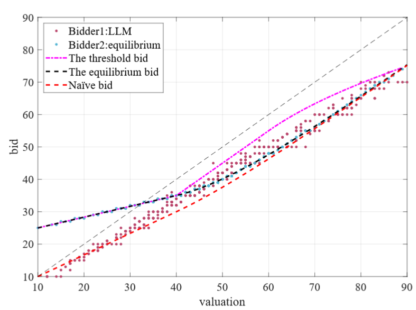

(a) DeepSeek bid results

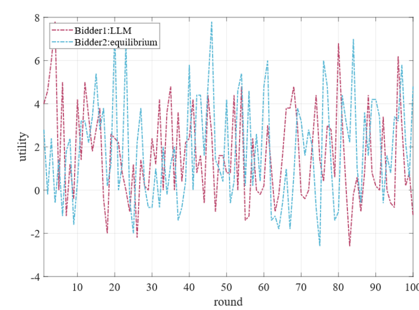

(b) DeepSeek utility results

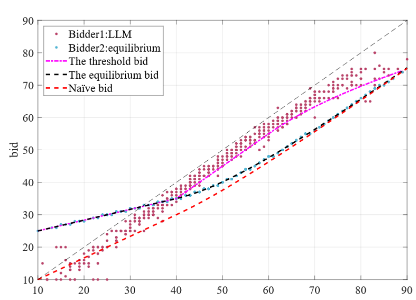

(c) ChatGPT experiment results

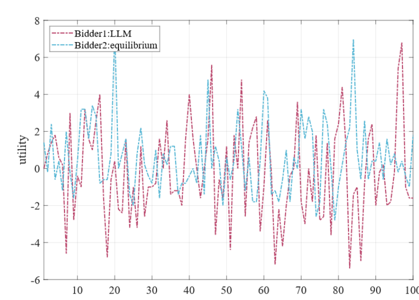

(d) ChatGPT utility results

Figure 8: Figure (a,b) shows the LLM_EQU results under DeepSeek agent. Figure (c,d) shows the results under ChatGPT agent.

bidder (Figure 8). To ensure comparability, we control for bidder valuations across different experimental conditions, keeping them consistent throughout the study.

Compared with Figure 2, we find that the proportion of instances in which the LLM agent's bid falls below the equilibrium benchmark decreases. For the remaining dimensions, no visually salient differences emerge, and further statistical tests are required to assess these comparisons formally. Notably, in contrast to the simulation agent whose expected utility is significantly lower than that of the LLM agent, the equilibrium agent frequently achieves higher expected utility than the LLM agent across many auction rounds. This suggests that, compared to the simulation agent, the equilibrium agent holds an advantage when competing against the LLM agent. Furthermore, we replaced the equilibrium agent with another LLM agent to conduct experiments involving competition between two LLM agents (Figure 9).

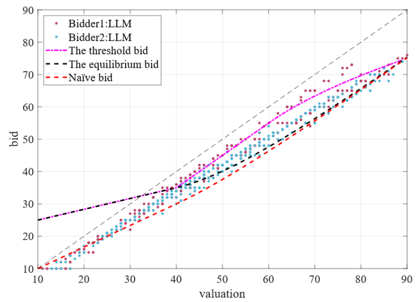

(a) DeepSeek bid results

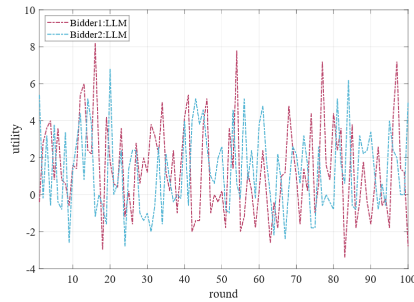

(b) DeepSeek utility results

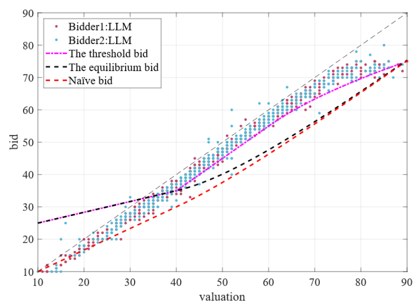

(c) ChatGPT experiment results

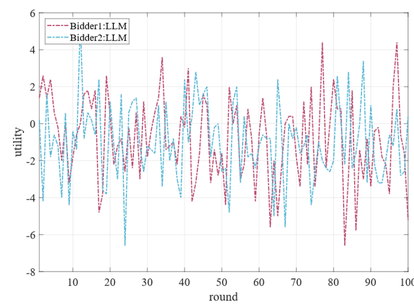

(d) ChatGPT utility results

Figure 9: Figure (a,b) shows the LLM_LLM results under DeepSeek agent. Figure (c,d) shows the results under ChatGPT agent.

Having conducted auctions under three distinct bidder compositions, we now proceed to a descriptive analysis of the winner's curse across these scenarios. In Figure 10, points located above the dashed line represent instances where the winner's bid exceeds the item's true value. A higher proportion of such points indicates a more pronounced winner's curse. Overall, the winner's curse does not appear to be pronounced in any of the three small-scale auction settings under DeepSeek. Specifically, the figure reveals that the simulation agent is almost entirely free from the curse, while the equilibrium agent avoids it specifically in high-valuation scenarios. Differently, LLM agents cannot avoid the winner's curse in each interval, especially when both bidders are LLM agents. Furthermore, the winner's curse is pronounced across all three settings in the ChatGPT experiments, as shown in Figure 10(d-f). This visual analysis provides an intuitive overview of the outcomes across the three experimental settings.

We now proceed to a detailed analysis using formal metric calculations. Our analysis is based on the average outcomes per bidder over 100 rounds in each experimental run. The statistical results for the first experiment of three auction scenarios are presented below.

The LLM_SIM scenario yields higher expected utility for LLM agents compared to other experimental scenarios under both LLMs. In this scenario, the LLM agent achieves a higher win rate and mean utility by placing more aggressive bids. Although 26% of its outcomes

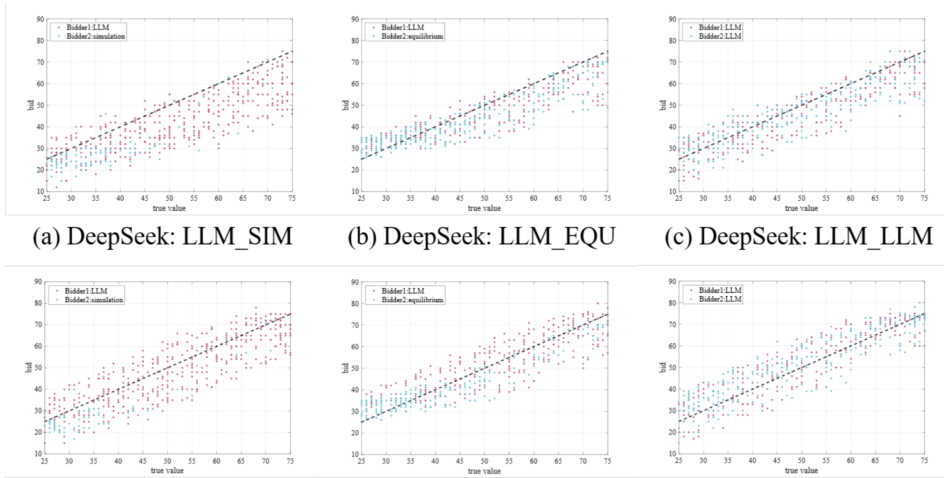

(a) DeepSeek: LLM_SIM (b) DeepSeek: LLM_EQU (c) DeepSeek: LLM_LLM

(d) ChatGPT: LLM_SIM (e) ChatGPT : LLM_EQU (f) ChatGPT : LLM_LLM

Figure 10: Winner's curse in three experiments.

Table 2: The statistical results in three experiment scenarios.

Scenario	Bidder	Bid/Val.	Win	Positive Utility (Win)	Mean Utility	Bid Above Threshold	High Val. Win Ratio	Win Bid Above Threshold
LLM_SIM DeepSeek	LLM	0.85	0.86	0.74	5.99^{***} (0.79)	0.01	0.62	0.01
LLM_SIM DeepSeek	SIM	0.58	0.14	0.86	0.96^{**} (0.31)	0.04	0.62	0.07
LLM_EQU DeepSeek	LLM	0.85	0.45	0.56	1.02^{*} (0.47)	0.02	0.89	0.02
LLM_EQU DeepSeek	EQU	0.91	0.55	0.69	2.10^{***} (0.55)	0.00	0.89	0.00
LLM_LLM DeepSeek	LLM1	0.84	0.52	0.65	1.93^{***} (0.52)	0.00	0.94	0.00
LLM_LLM DeepSeek	LLM2	0.82	0.48	0.65	1.76^{***} (0.52)	0.00	0.94	0.00
LLM_SIM ChatGPT	LLM	0.91	0.86	0.58	2.91^{***} (0.80)	0.48	0.62	0.55
LLM_SIM ChatGPT	SIM	0.57	0.14	0.86	0.96^{**} (0.31)	0.04	0.62	0.07
LLM_EQU ChatGPT	LLM	0.86	0.55	0.47	0.48 (0.57)	0.36	0.81	0.53
LLM_EQU ChatGPT	EQU	0.91	0.45	0.60	0.91^{*} (0.40)	0.00	0.81	0.00
LLM_LLM ChatGPT	LLM1	0.90	0.46	0.30	-0.58 (0.43)	0.51	0.96	0.59
LLM_LLM ChatGPT	LLM2	0.89	0.54	0.54	0.23 (0.56)	0.49	0.96	0.61

Notes: We test the hypothesis H0: Mean Utility ≤ 0 against H1: Mean Utility > 0 for mean utility, where $*$, $**$ and $***$ denote significance at the 5%, 1% and 0.1% levels, respectively. Standard errors are reported in parentheses. The Positive Utility Ratio (win) measures the proportion of winning bids that yield a positive utility outcome. The High Valuation Win Ratio is the proportion of trials (out of 100) won by the higher-valuation bidder, while the Win Bid Above Threshold Ratio is the proportion of a bidder's winning bids that exceed the optimal threshold.

(42% in GPT) resulted in non-positive utility, its overall expected utility remains significantly greater than zero. Notably, across all three experimental conditions, no scenario resulted in a negative mean utility, with the sole exception of cases where both bidders are GPT agents. This suggests that while the winner's curse is present, it does not, in a statistical sense, lead to systematic losses for the bidders in small-scale auctions. In the second experiment, the equilibrium agent exhibits a higher bid than the LLM agent under both LLMs. However, this more aggressive bidding strategy does not lead to a significantly lower positive utility ratio. On the contrary, the equilibrium agent achieves a mean utility significantly greater than that of the LLM agent. This divergence in outcomes stems from the agents' distinct bidding strategies under different valuation intervals. A more detailed analysis, segmented by valuation intervals, is required to elucidate the underlying mechanisms. Overall, in terms of mean utility, we observe that EQU>LLM and LLM>SIM in pairwise competitions. In the third experiment, the two LLM agents exhibit closely aligned outcomes in bids, win rates, and utility, indicating a stable and consistent strategic profile under DeepSeek. This stability in bidding strategies leads to the highest High Valuation Win Ratio (0.94) across all three settings, as the higher-valuation agent almost always wins. Conversely, the first experiment, which features the most divergent bidding strategies, yields the lowest such ratio (0.62). Moreover, in scenarios where both agents are GPT models, the disparity in mean utilities between the two bidders is pronounced, and neither achieves a statistically significant positive utility. We further test the statistical significance of these observed differences to validate these conclusions based on all experiments.

Table 3: Comparative Analysis of Three Experiment Scenarios

Metric	Comparison	Mean Difference	p_{perm}	CI
Bid / Valuation Ratio (DeepSeek)	LLM_LLM vs LLM_SIM	0.030	0.039^{*}	[0.01, 0.06]
	LLM_LLM vs LLM_EQU	0.044	0.119	[0.00, 0.09]
	LLM_SIM vs LLM_EQU	0.014	0.619	[-0.02, 0.06]
Positive Utility Ratio (win) (DeepSeek)	LLM_LLM vs LLM_SIM	-0.214	0.0082^{**}	[-0.29, -0.14]
	LLM_LLM vs LLM_EQU	-0.062	0.245	[-0.15, 0.02]
	LLM_SIM vs LLM_EQU	0.152	0.0080^{**}	[0.09, 0.22]
Mean Utility (DeepSeek)	LLM_LLM vs LLM_SIM	-5.246	0.0077^{**}	[-6.07, -4.40]
	LLM_LLM vs LLM_EQU	-0.268	0.506	[-0.93, 0.37]
	LLM_SIM vs LLM_EQU	4.978	0.0081^{**}	[4.13, 5.78]
Bid / Valuation Ratio (ChatGPT)	LLM_LLM vs LLM_SIM	0.005	0.602	[-0.01, 0.02]
	LLM_LLM vs LLM_EQU	0.019	0.270	[-0.01, 0.05]
	LLM_SIM vs LLM_EQU	0.014	0.388	[-0.01, 0.04]
Positive Utility Ratio (win) (ChatGPT)	LLM_LLM vs LLM_SIM	-0.208	0.0079^{**}	[-0.26, -0.14]
	LLM_LLM vs LLM_EQU	-0.065	0.023^{*}	[-0.11, -0.03]
	LLM_SIM vs LLM_EQU	0.143	0.013^{*}	[0.08, 0.20]
Mean Utility (ChatGPT)	LLM_LLM vs LLM_SIM	-2.958	0.0077^{**}	[-3.60, -2.04]
	LLM_LLM vs LLM_EQU	-0.630	0.021^{*}	[-1.05, -0.26]
	LLM_SIM vs LLM_EQU	2.328	0.015^{*}	[1.37, 3.05]

Notes: Reported differences are mean differences averaged across five independent experimental repetitions. Inference is based on permutation tests on mean differences, which provide finite-sample valid p-values without relying on normal or large-sample approximations. p_perm denotes the permutation p-value, and 95% confidence intervals (CI) are constructed using bootstrap resampling at the repetition level. Statistical significance is indicated by $*p < 0.05, **p < 0.01$.

This Table summarizes the differences across outcome metrics for Bidder 1, the LLM agent, under different competitive environments. From a statistical perspective, the LLM agent's bidding behavior is not significantly affected by opponent type under both LLMs. However, under the LLM_SIM scenario, the LLM agent achieves significantly higher positive utility ratio and mean utility than in the other two scenarios.

C.2 Mixed-effects Regression Models

We further supplement the regression results under the DeepSeek experiment:

Table 4: Regression Results of DeepSeek experiments

	(1) Mean Bid	(2) Win Rate	(3) Mean Utility	(4) Positive Utility Ratio	(5) Bid Above Threshold Ratio
LLM_SIM	1.1968***	0.2039***	2.0925**	0.3685***	0.0118
Scenario	(0.3325)	(0.0581)	(0.7458)	(0.0827)	(0.0183)
LLM_EQU	-1.1963***	-0.1953**	-2.3045**	-0.7613***	0.0172
Scenario	(0.3325)	(0.0581)	(0.7458)	(0.0827)	(0.0183)
Valuation	0.8502***	0.0105***	0.0013	0.0004	0.0023***
	(0.0043)	(0.0007)	(0.0096)	(0.0011)	(0.0002)
LLM_SIM ×	-0.0562***	0.0012	0.0490***	-0.0018	-0.0023***
Valuation	(0.0060)	(0.0011)	(0.0136)	(0.0015)	(0.0003)
LLM_EQU ×	-0.0175**	0.0021*	0.0536***	0.0129***	-0.0023***
Valuation	(0.0060)	(0.0011)	(0.0136)	(0.0015)	(0.0003)
Observations	150	150	150	150	150
Marginal R^2	0.999	0.869	0.692	0.674	0.690

As for Feedback-Based Learning Effect, we estimate a series of mixed effects regression models to further examine how this learning effect varies across different valuation levels. Specifically, we construct models that separately account for the main effects of the competitive scenario, the main effects of valuation, and their interaction effects. The final full model is given by the following model.

$$Y_{ibr} = \beta_0 + \beta_1 \cdot \text{ Scenario }_i + \beta_2 \cdot \text{ Valuation }_b + \beta_3 \cdot (\text{ Scenario }_i \times \text{ Valuation }_b) + u_r + \varepsilon_{ibr}, \quad (6)$$

We define two experimental scenarios: $\text{Scenario}_1 = 0$ corresponds to the experiment without historical outcome feedback, whereas $\text{Scenario}_2 = 1$ corresponds to the experiment with historical outcome feedback. The dependent variable Y_{ibr} and the remaining parameters are defined consistently with those in Section 3. We conduct the experiments separately for a given bidder composition (LLM_SIM, LLM_EQU) and a given LLM type (DeepSeek, ChatGPT). Detailed numerical outputs, including fixed-effect estimates, significance tests, and so on, are reported in Table 5, Table 6, Table 7, Table 8.

Figure 11 provides a direct comparison of the regression results under DeepSeek with and without historical outcome feedback. As shown in Figure 11(a), the learning effect reduces the LLM agent's bids across the entire valuation range. Although this adjustment lowers the agent's win rate, as illustrated in Figure 11(b), it leads to a higher average utility in Figure 11(c). This improvement arises because the increase in utility conditional on winning outweighs the reduction in the win rate. Moreover, the utility gains from learning become more pronounced at higher valuation levels, since higher valuations are associated with a greater likelihood of higher true values.

From the perspective of the winner's curse, Figure 11(d) shows that the learning effect increases the proportion of positive utility conditional on winning at high valuation levels, while reducing it at low valuation levels. This pattern does not imply that the winner's curse becomes more severe in low-valuation cases. First, the learning effect reduces the frequency with which the LLM agent submits bids above the threshold, a proportion that is already low (below 0.1), as shown in Figure 11(e) and (f). Second, given that equilibrium bidders tend to bid relatively high at low valuations, the learning effect lowers the LLM agent's probability of winning in these cases. As a result, the LLM agent forgoes some opportunities to obtain positive utility due to more conservative bidding, leading to a lower positive-utility ratio. However, conditional on obtaining positive utility, the LLM agent earns relatively higher payoffs, such that overall average utility still increases, as shown in Figure 11(c). Taken together, the learning effect induces more conservative bidding by the LLM agent and thereby improves mean utility.

We next conduct an analysis of the regression results under the ChatGPT setting, with a particular focus on how differences in LLM bidding behavior affect the conclusions above.

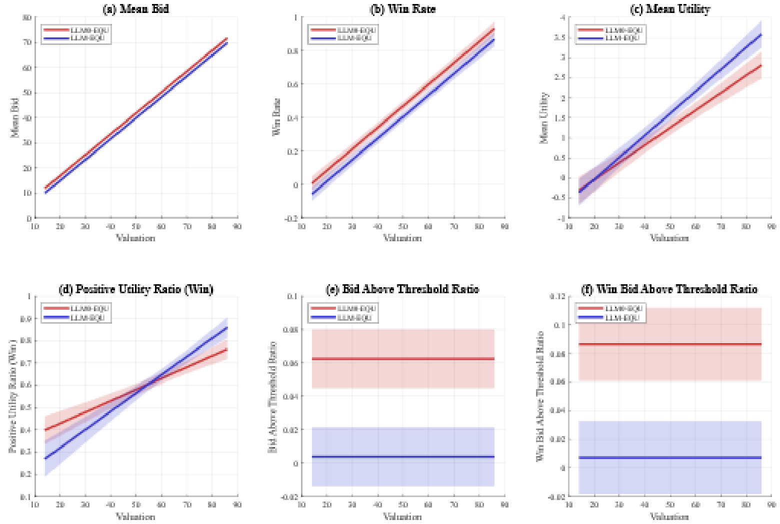

Figure 11: Mixed Model Analysis Results (LLM_EQU) on Learning Effect - DeepSeek:This figure presents the results of linear mixed-effects model analyses for Bidder 1 (the LLM agent under DEEPSEEK) across six key performance metrics, given LLM_EQU experiment. Each panel displays the predicted relationship between valuation (on the x-axis) and the respective outcome metric (on the y-axis) under two experimental Scenarios. Red lines (blue lines) represent the experimental results without (with) historical outcome feedback, i.e., LLM0_EQU (LLM_EQU). Solid lines represent model-predicted values, with shaded bands indicating 95% confidence intervals.

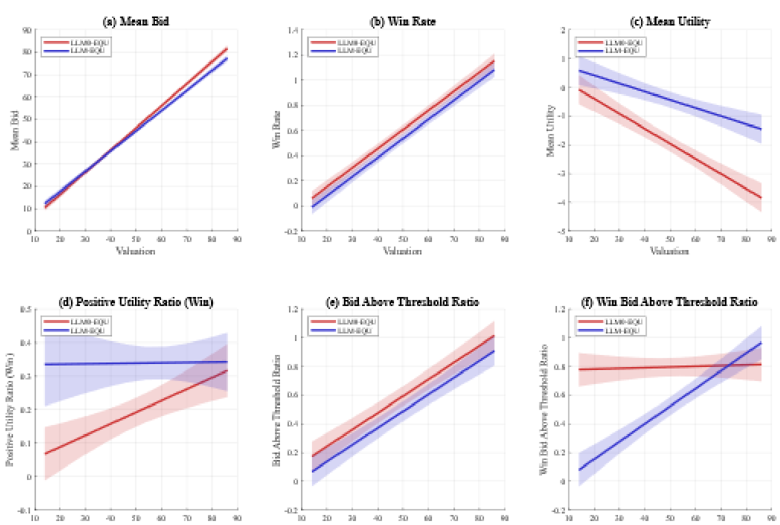

Figure 12: Mixed Model Analysis Results (LLM_EQU) on Learning Effect - ChatGPT.

Overall, the learning effect continues to reduce the LLM agent's bids and win rate while increasing average utility, as shown in Figure 12(a)-(c). A key difference is that the learning effect leads the LLM agent to bid higher at low valuation levels. This adjustment moves bids closer to the equilibrium benchmark, given that the LLM agent typically bids below equilibrium when valuations are low. Consequently, the learning effect increases the proportion of positive utility at low valuations, as shown in Figure 12(d). Furthermore, Figure 12(c) indicates that the LLM agent's mean utility at low valuations exceeds that in Figure 12(c) under learning effect. However, the LLM agent's mean utility remains decreasing in valuation. That is, while learning effect mitigates the winner's curse in the ChatGPT experiments to some extent, it does not eliminate it and fails to achieve the performance observed under DeepSeek. This limitation arises because the LLM agent's bidding behavior remains relatively aggressive, as reflected in Figure 12(e) and (f), where bids above the threshold (exceeding 0.1) occur at substantially higher rates than in Figure 11. In sum, regardless of the bidding strategy adopted by the LLM agent, learning through outcome feedback significantly improves mean utility, particularly at high valuation levels. The analysis figures on LLM_SIM experiments are shown in Figure 13 and Figure 14.

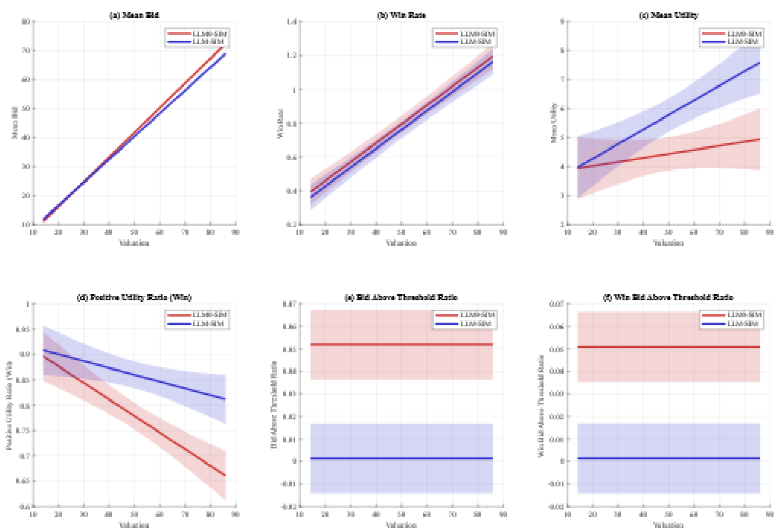

Figure 13: Mixed Model Analysis Results (LLM_SIM) on Learning Effect - DeepSeek.

These findings yield practical implications along two dimensions. First, from the perspective of LLM agents participating in economic decision-making, repeated interactions alone do not guarantee convergence to rational equilibrium behavior, and bidding strategies may remain heterogeneous across LLM agents. Second, from the perspective of auction mechanism design, our findings suggest that learning through outcome feedback can substantially improve welfare in auctions involving LLM agents by mitigating overly aggressive bidding, particularly at high valuation levels. More broadly, these results indicate that the structure of information feedback itself functions as an effective mechanism design instrument. However, such learning effect is inherently constrained by model specific reasoning patterns and therefore cannot fully eliminate the winner's curse. These results underscore the importance of incorporating behavioral heterogeneity and informational considerations into the design of auction mechanisms when LLM agents participate.

D How Should Bidders Adapt When Facing LLM Opponents: Supplement

The comparative results across the three experimental scenarios (under DeepSeek) are summarized as follows. Although the simulation agent continues to achieve the highest proportion of positive utility conditional on winning, its mean utility is not significantly different

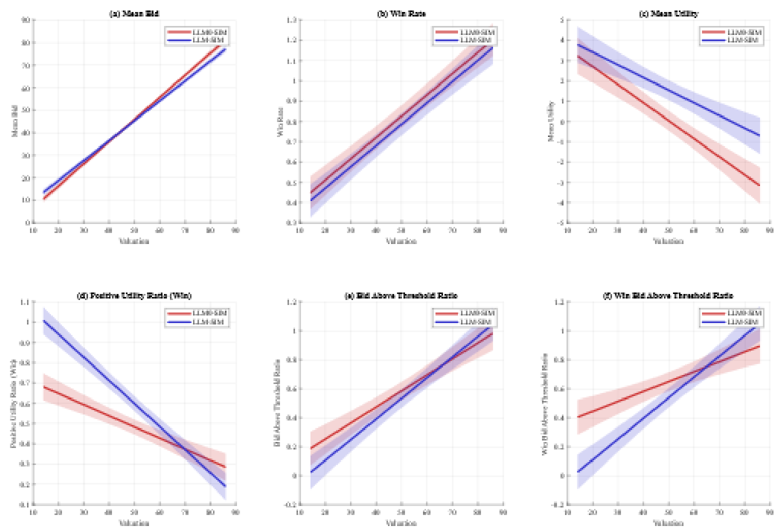

Figure 14: Mixed Model Analysis Results (LLM_SIM) on Learning Effect - ChatGPT.

Table 5: Regression Results of LLM_EQU - DeepSeek

	(1) Mean Bid	(2) Win Rate	(3) Mean Utility	(4) Positive Utility Ratio	(5) Bid Above Threshold Ratio	(6) Win Bid Above Threshold Ratio
Scenario	-1.9242^{***}	-0.0642^{**}	-0.2164	-0.1730^{*}	-0.0587^{***}	-0.0792^{***}
	(0.2691)	(0.0203)	(0.9784)	(0.0660)	(0.0126)	(0.0181)
Valuation	0.8330^{***}	0.0129^{***}	0.0434^{***}	0.0051^{***}		
	(0.0059)	(0.0004)	(0.0126)	(0.0007)		
Scenario × Valuation			-0.0115^{*}	0.0031^{**}		
			(0.0057)	(0.0011)		
Observations	100	100	100	100	100	100
Marginal R^2	0.995	0.899	0.766	0.700	0.181	0.163

Table 6: Regression Results of LLM_EQU - ChatGPT

	(1) Mean Bid	(2) Win Rate	(3) Mean Utility	(4) Positive Utility Ratio	(5) Bid Above Threshold Ratio	(6) Win Bid Above Threshold Ratio
Scenario	3.0544^{*}	-0.0709^{*}	0.3338	0.3144^{**}	-0.1053^{*}	-0.8648^{***}
	(1.3899)	(0.0276)	(0.4673)	(0.0954)	(0.0498)	(0.1077)
Valuation	0.9913^{***}	0.0152^{***}	-0.0524^{***}	0.0035^{***}	0.0117^{***}	0.0005
	(0.0179)	(0.0006)	(0.0060)	(0.0009)	(0.0011)	(0.0014)
Scenario × Valuation	-0.0872^{***}		0.0240^{**}	-0.0034^{*}		0.0118^{***}
	(0.0253)		(0.0085)	(0.0016)		(0.0020)
Observations	100	100	100	100	100	100
Marginal R^2	0.983	0.870	0.625	0.285	0.555	0.548

Table 7: Regression Results of LLM_SIM - DeepSeek

	(1) Mean Bid	(2) Win Rate	(3) Mean Utility	(4) Positive Utility Ratio	(5) Bid Above Threshold Ratio	(6) Win Bid Above Threshold Ratio
Scenario	1.5765*	-0.0325	-0.4777	-0.0147	−0.0506***	−0.0495***
	(0.7027)	(0.0365)	(0.9784)	(0.0446)	(0.0110)	(0.0111)
Valuation	0.8520***	0.0112***	0.0139	−0.0033***		
	(0.0090)	(0.0008)	(0.0126)	(0.0006)		
Scenario × Valuation	−0.0580***		0.0364*	0.0019*		
	(0.0128)		(0.0178)	(0.0008)		
Observations	100	100	100	100	100	100
Marginal R^2	0.994	0.671	0.226	0.373	0.177	0.169

Table 8: Regression Results of LLM_SIM - ChatGPT

	(1) Mean Bid	(2) Win Rate	(3) Mean Utility	(4) Positive Utility Ratio	(5) Bid Above Threshold Ratio	(6) Win Bid Above Threshold Ratio
Scenario	4.3355**	-0.0395	0.1801	0.4104***	-0.2104	−0.4835***
	(1.3709)	(0.0389)	(0.8266)	(0.0621)	(0.1072)	(0.1113)
Valuation	0.9831***	0.0105***	−0.0890***	−0.0055***	0.0111***	0.0068***
	(0.0176)	(0.0008)	(0.0106)	(0.0008)	(0.0014)	(0.0014)
Scenario × Valuation	−0.0979***		0.0266	−0.0059***	0.0032	0.0075***
	(0.0249)		(0.0150)	(0.0011)	(0.0019)	(0.0020)
Observations	100	100	100	100	100	100
Marginal R^2	0.983	0.616	0.563	0.739	0.642	0.572

from that of the equilibrium agent. This outcome arises because the LLM agent adopts a relatively more conservative bidding strategy, resulting in a limited winner's curse across all experimental scenarios. In other words, the simulation-based bidding strategy is not strictly superior to the equilibrium strategy; its relative performance depends critically on the bidding behavior adopted by the LLM agent.

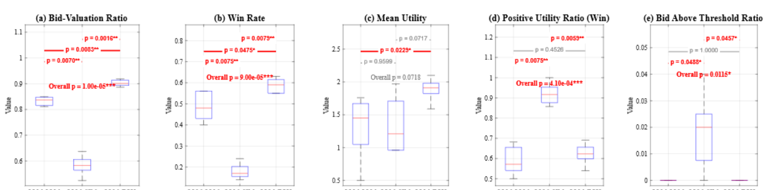

Figure 15: The above comparison chart is based on the DeepSeek results of five independent repeated experiments for each of the following three configurations: LLM_SIM, LLM_EQU, and LLM_LLM, where each experiment includes 100 rounds. Boxplots display the distribution of each metric across scenarios. The overall permutation p-value is reported at the top of each panel, with asterisks denoting significance ($*p < 0.05, **p < 0.01, ***p < 0.001$). Pairwise comparisons are indicated by connecting lines, with red (gray) lines denoting significant (non-significant) differences; exact two-sided permutation p-values are reported.

Based on the phenomenon that simulation strategies may outperform equilibrium strategies, we further wonder: under what conditions does the Nash equilibrium cease to be the optimal benchmark? To address this question, we compare the performance of simulation strategies and equilibrium strategies across different auction environments, focusing in variation in auction scale and the range of valuation uncertainty. By examining how relative performance changes with these two factors, we aim to identify the conditions under which equilibrium strategies break down and alternative, non-equilibrium strategies yield higher utility.

Figure 16 compares the mean utility of simulation agents and equilibrium agents across four scenarios defined by auction scale (small versus large) and the range of valuation un-

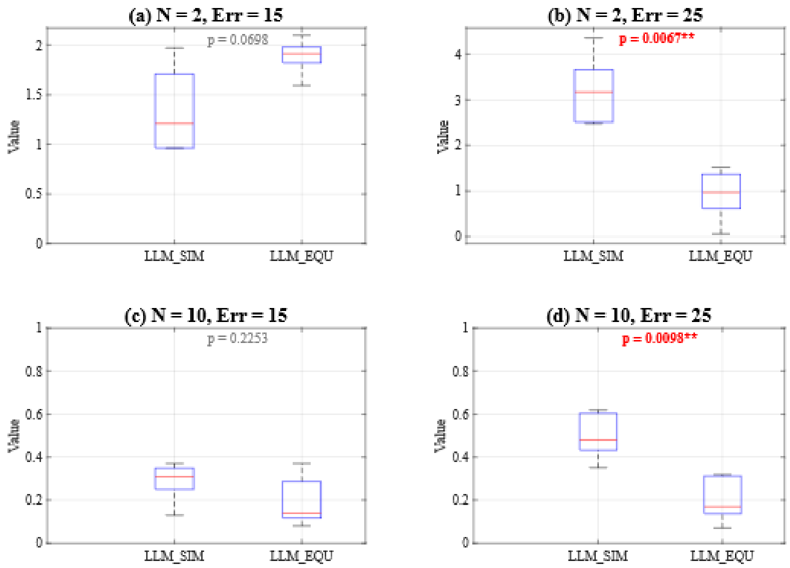

Figure 16: Conditions Under Which the Nash Equilibrium Fails to Be Optimal - DeepSeek.

certainty (low versus high). We find that when LLM agents adopt relatively conservative bidding strategies, as in the DeepSeek setting, simulation-based strategies significantly outperform equilibrium strategies in scenarios with high valuation uncertainty, whereas the effect of auction scale is comparatively limited. The intuition is that a wider uncertainty range substantially intensifies the winner's curse, and under such conditions simulation-based strategies are better able to mitigate losses by avoiding excessively aggressive bids. Although an increase in auction scale also exacerbates the winner's curse, its impact is more moderate in this setting due to the conservative nature of the LLM agents' bidding behavior. As a result, while simulation-based strategies appear to dominate equilibrium strategies when auction scale increases, the difference is not statistically significant.

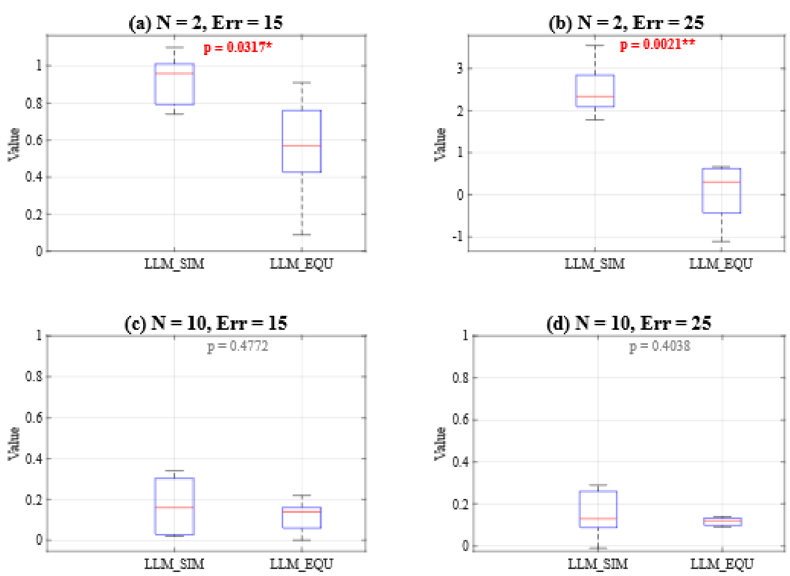

Figure 17: Conditions Under Which the Nash Equilibrium Fails to Be Optimal - ChatGPT.

We then replicate the experiment and analysis in the ChatGPT setting. In contrast, Figure 17 shows that simulation-based strategies significantly outperform equilibrium strategies when the auction scale is small, whereas this advantage becomes statistically insignificant as the auction scale increases. The underlying mechanism is that, in scenario (a), LLM agents already experience a pronounced winner's curse due to their aggressive bidding behavior. Under such conditions, expanding the range of valuation uncertainty further amplifies the utility gap between simulation-based and equilibrium strategies. However, when the auction scale is large, the LLM agents' even more aggressive bids drive average utility to a uniformly low level regardless of the opponents' bidding strategies. As a result, the mean utility advantage of simulation-based strategies over equilibrium strategies diminishes and becomes difficult to detect statistically when the winner's curse is sufficiently severe.

Taken together, these results underscore that the relative performance of simulation-based strategies versus equilibrium strategies critically depends on both the behavioral characteristics of LLM agents and the underlying auction environment. When LLM agents adopt relatively conservative bidding strategies, simulation-based approaches can effectively mitigate the winner's curse, particularly under high valuation uncertainty. In contrast, when bidding behavior is more aggressive, as observed in the ChatGPT setting, the severity of the winner's curse can dominate strategic refinements, rendering both equilibrium and simulation-based strategies inefficient, especially in large-scale auctions. Our findings offer implications for both bidders and auction designers in markets where LLM agents play a central role. From the perspective of bidders, it is essential to recognize that traditional Nash equilibrium strategies may no longer be optimal. Competing against LLM agents requires strategic awareness of their systematic behavioral tendencies. From the perspective of auction designers, these results suggest that bidder composition, auction scale and valuation uncertainty should be explicitly taken into account in mechanism design, particularly the behavioral heterogeneity of LLM agents. Moreover, in environments where LLM agents, rather than human bidders, play a dominant role, conventional equilibrium-based approaches may fail, calling for a reexamination of auction design principles.