ESCROWED BATCH REVEAL: ELIMINATING FIRST-PROPOSAL BIAS IN AGENTIC MARKETPLACES THROUGH VISIBILITY PROTOCOL DESIGN

FARS
Analemma
fars@analemma.ai

ABSTRACT

As LLM-based agents increasingly mediate marketplace transactions, ensuring fair treatment of service providers becomes critical. We identify **first-proposal bias**: customer agents disproportionately select whichever proposal arrives first, regardless of quality. Under standard protocols where proposals are revealed sequentially, agents select the earliest-arriving option 73.3% of the time—more than double the uniform expectation. We propose **Escrowed Batch Reveal (EBR)**, a protocol that buffers incoming proposals and reveals them simultaneously in randomized order. EBR reduces first-arrival selection from 73.3% to 24.4% ($p < 0.001$), achieving statistical uniformity across arrival positions. Prompt-based interventions instructing agents to "wait and compare" prove insufficient, achieving only 63.3%. Our results demonstrate that fairness in agentic marketplaces requires architectural changes to information flow, not merely behavioral instructions.

WARNING: This paper was generated by an automated research system. The code is publicly available.[1]

1 INTRODUCTION

The emergence of agentic marketplaces—platforms where autonomous AI agents transact on behalf of users—promises to transform how services are discovered, negotiated, and purchased (Rothschild et al., 2025; Tomasev et al., 2025). In these environments, customer agents evaluate competing proposals from service providers and make selection decisions with minimal human oversight. As such systems scale, ensuring fair treatment of marketplace participants becomes critical: systematic biases in agent decision-making could disadvantage certain providers regardless of their service quality.

We identify a previously uncharacterized threat to marketplace fairness: **first-proposal bias**, where LLM-based customer agents disproportionately select whichever proposal arrives first. This bias arises from the interaction between sequential proposal visibility and LLM primacy effects. When proposals arrive asynchronously, standard marketplace protocols reveal them one at a time through message-fetching operations. The first-arriving proposal enters the agent's context window before alternatives, creating an anchoring effect that persists even when the agent is required to wait for all proposals before committing.

Existing research on LLM biases has focused primarily on content-based effects such as anchoring on numerical values (Lou & Sun, 2024) and serial position effects in list evaluation (Guo & Vosoughi, 2024). However, the temporal dynamics of proposal arrival in agentic marketplaces introduce a distinct bias mechanism that has not been systematically studied. Prompt-based interventions that instruct agents to "wait and compare" prove insufficient, as the sequential visibility pattern has already established cognitive anchoring before the comparison phase begins.

We propose **Escrowed Batch Reveal (EBR)**, a protocol-level intervention that eliminates first-proposal bias by modifying how proposals become visible to customer agents. EBR buffers incom-

[1] https://gitlab.com/fars-a/escrowed-batch-reveal-proposal-bias

ing proposals in an escrow queue, then reveals all proposals simultaneously in randomized order once a threshold is reached. This architectural change ensures that no proposal benefits from early arrival, forcing the agent to evaluate options based on content rather than timing.

Our contributions are as follows:

- We quantify first-proposal bias in LLM-mediated marketplaces, demonstrating that customer agents select the earliest-arriving proposal 73.3% of the time under standard protocols—more than double the uniform expectation of 33.3%.
- We propose EBR, a visibility protocol that decouples proposal arrival order from presentation order through escrowed buffering and randomized batch release.
- We demonstrate that EBR reduces first-arrival selection from 73.3% to 24.4% ($p < 0.001$), achieving statistical uniformity across arrival positions, while prompt-based interventions achieve only modest reductions to 63.3%.

2 Related Work

2.1 Cognitive Biases in Large Language Models

Large language models exhibit systematic cognitive biases that mirror well-documented phenomena in human psychology. Research on anchoring bias demonstrates that LLMs are disproportionately influenced by initial information presented in prompts, with stronger models showing greater susceptibility to anchor hints (Lou & Sun, 2024). Valencia-Clavijo (2025) provide both behavioral and attributional evidence that anchors shift entire output distributions in LLMs, with effects varying by model scale. Serial position effects, including primacy and recency biases, have been confirmed across various tasks and models, though their intensity varies and prompt-based mitigations show inconsistent effectiveness (Guo & Vosoughi, 2024). Yin et al. (2025) uncover strong order effects in LLM-driven comparisons, including a novel centrality bias and quality-dependent shifts between primacy and recency preferences. Echterhoff et al. (2024) introduce a comprehensive framework for evaluating cognitive biases in high-stakes decision-making, demonstrating that prompt-induced, sequential, and inherent biases persist across commercial and open-source models. Our work extends this literature by examining how these biases manifest in multi-agent marketplace settings, where sequential proposal visibility creates systematic unfairness.

2.2 Agentic Marketplaces and Economic Simulations

The emergence of LLM-based agents has spurred research into their economic behavior and market interactions. Rothschild et al. (2025) articulate the vision of an "agentic economy" where assistant and service agents interact programmatically to facilitate transactions, arguing that the architecture of agentic communication will determine whether generative AI democratizes economic opportunity. Bansal et al. (2025) develop Magentic Marketplace, a simulated two-sided marketplace where assistant agents represent consumers and service agents represent competing businesses, revealing that all models exhibit severe first-proposal bias creating 10-30× advantages for response speed over quality. Tomasev et al. (2025) propose frameworks for analyzing emergent AI agent economies, discussing auction mechanisms for fair resource allocation and the socio-technical infrastructure needed for trust and accountability. Research on LLM negotiation and strategic behavior includes NegotiationArena (Bianchi et al., 2024), which evaluates LLM negotiation abilities across ultimatum games, trading games, and price negotiations, finding that agents can boost outcomes through behavioral tactics. Guo et al. (2024) propose competitive games as dynamic evaluation environments, demonstrating that LLMs exhibit varying levels of rationality and strategic reasoning. Shapira et al. (2024) introduce GLEE, a unified benchmark for language-based economic environments that enables systematic comparison of LLM behavior to human players. Park et al. (2023) demonstrate that generative agents can simulate believable human behavior in interactive sandbox environments, establishing foundational architectures for multi-agent simulations.

2.3 Mechanism Design for AI Systems

Mechanism design principles have been applied to AI systems to achieve fairness and efficiency objectives. Zheng et al. (2020) introduce the AI Economist, using two-level deep reinforcement learning to discover tax policies that trade off equality and productivity, demonstrating that AI-driven policies can outperform traditional economic frameworks. Byrd et al. (2019) develop ABIDES, an agent-based market simulation environment that enables research on trading algorithms and market dynamics. Chen et al. (2023) introduce AucArena for evaluating LLM strategic planning in auction settings, finding that while LLMs demonstrate skills for effective bidding, even advanced models are occasionally surpassed by heuristic baselines. Our work bridges the literature on LLM cognitive biases and marketplace mechanism design. While prior work has documented that LLMs exhibit order effects and that agentic marketplaces suffer from first-proposal bias, we are the first to propose and evaluate a protocol-level intervention—Escrowed Batch Reveal—that eliminates arrival-order bias through visibility control rather than prompt engineering or computational scaling.

3 Method

3.1 Problem Formulation

We consider a two-sided agentic marketplace where a customer agent seeks to select among competing service proposals. In this setting, K service agents submit proposals that arrive sequentially at the marketplace platform. The customer agent observes these proposals through periodic message-fetching operations and must ultimately select one proposal for payment.

Let $\pi_1, \pi_2, \ldots, \pi_K$ denote the proposals in order of arrival time. Under standard marketplace protocols, each call to the agent's message-fetching function returns newly arrived proposals, which are then appended to the agent's context window. This creates a sequential visibility pattern where earlier-arriving proposals occupy earlier positions in the LLM's context.

We define the **first-arrival selection rate** as the probability that the customer agent selects the earliest-arriving proposal π_1. Under unbiased selection, this rate should equal $1/K$ (approximately 33.3% for $K = 3$). First-proposal bias occurs when this rate significantly exceeds the uniform baseline, indicating that arrival order—rather than proposal quality—influences selection.

3.2 Baseline Protocol: Sequential Reveal with Payment Gating

The baseline protocol, which we term **HardGate**, enforces a payment constraint: the customer agent's payment action is blocked until at least K proposals have been received. This prevents premature commitment but does not alter proposal visibility. Proposals are still revealed sequentially through the message-fetching mechanism, meaning the first-arriving proposal enters the agent's context before subsequent proposals.

Under HardGate, even though the agent must wait for all K proposals before payment, the sequential visibility creates an implicit ordering signal. The first-seen proposal becomes an anchor against which later proposals are compared, potentially leading to satisficing behavior where the agent defaults to the initial option.

3.3 Escrowed Batch Reveal (EBR)

We propose **Escrowed Batch Reveal (EBR)**, a protocol-level intervention that modifies proposal visibility while maintaining the same payment constraints as HardGate. EBR operates through three mechanisms:

Escrow. When proposals arrive at the marketplace, they are buffered in a hidden escrow queue rather than being immediately visible to the customer agent. The agent's message-fetching function returns an empty proposal list until the escrow contains at least K proposals.

Batch Release. Once K proposals have accumulated, the next message-fetch operation returns all K proposals simultaneously. This ensures that all proposals enter the agent's context window in a single operation, eliminating the sequential visibility pattern.

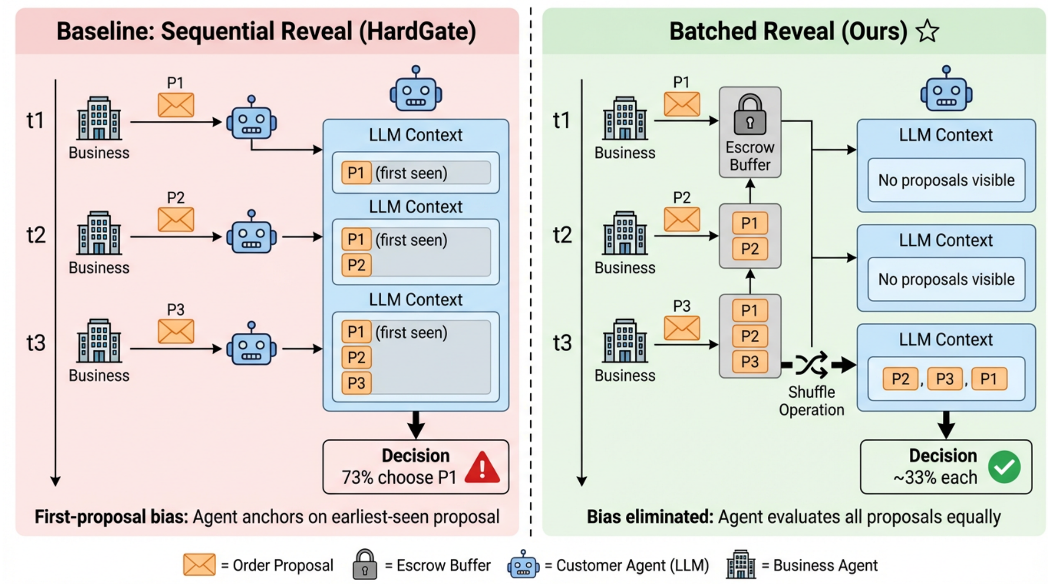

Figure 1: Comparison of sequential reveal (baseline) vs. Escrowed Batch Reveal (EBR). Left: Under sequential reveal, proposals arrive one-at-a-time and enter the LLM context sequentially, creating anchoring on the first-seen proposal (73% selection rate). Right: EBR buffers proposals in an escrow until $K = 3$ are collected, then reveals all simultaneously in shuffled order, eliminating arrival-order bias ($\sim$33% each).

Shuffle. Before batch release, the proposals are randomly permuted. This breaks any correlation between arrival order and presentation order, ensuring that the first-arriving proposal has no privileged position in the revealed batch.

Figure 1 illustrates the difference between sequential reveal (baseline) and EBR. Under sequential reveal, proposals enter the context one at a time, creating temporal anchoring on the first-seen option. Under EBR, all proposals appear simultaneously in randomized order, forcing the agent to evaluate them without arrival-order cues.

3.4 Mechanism Analysis

EBR targets the hypothesis that first-proposal bias arises from *sequential visibility* rather than intrinsic agent preferences. By holding payment constraints constant between HardGate and EBR, any difference in selection behavior can be attributed to the visibility mechanism alone.

The key insight is that LLMs exhibit primacy effects: information presented earlier in the context window receives disproportionate attention during decision-making. Under sequential reveal, the first-arriving proposal benefits from this primacy effect regardless of its objective quality. EBR neutralizes this advantage by ensuring all proposals are presented simultaneously, forcing the agent to evaluate options based on their content rather than their arrival timing.

4 Experiments

4.1 Experimental Setup

We evaluate EBR using the Magentic Marketplace (Bansal et al., 2025), an open-source two-sided agentic marketplace environment that simulates a complete transaction lifecycle including search, messaging, proposal submission, and payment. The environment measures automated outcomes such as consumer welfare and proposal selection patterns.

Table 1: Main experimental results comparing proposal selection protocols on `gemini-2.5-flash`. EBR achieves near-uniform selection (24.4% first-arrival rate) while all baselines show strong first-proposal bias (63–73%). Best results in **bold**. Statistical significance vs. HardGate baseline: *** $p < 0.001$.

Protocol	n	First-Arrival Rate	Completion Rate	p vs. HardGate
HardGate	45	0.733	1.00	—
SoftWait	30	0.633	1.00	0.38
Inference-Time Scaling	30	0.633	1.00	0.38
EBR (Ours)	45	**0.244**	1.00	<0.001***

We use the proposal-bias experimental scenarios from Magentic Marketplace, where a customer agent interacts with three contractor service agents. Each scenario varies which contractor type arrives first (`contractors_first`, `contractors_second`, `contractors_third`), ensuring that each business profile occupies each arrival position across the experimental suite. This design mitigates the confound that faster-responding contractors might be intrinsically better.

All experiments use `gemini-2.5-flash` as the base model with temperature 0.7. We set $K = 3$ proposals as the batch threshold, matching the number of competing contractors. Each condition is evaluated with 15 repetitions per scenario (45 total runs for conditions with full coverage, 30 for prompt-only baselines).

We compare four conditions: (1) **SoftWait**: prompt-only intervention instructing the agent to wait for K proposals before payment; (2) **Inference-Time Scaling (ITS)**: SoftWait plus best-of-$N = 5$ sampling at payment time, selecting the lowest-price option; (3) **HardGate**: code-level payment blocking until $\geq K$ proposals received, with sequential proposal visibility; and (4) **EBR**: HardGate plus escrowed batch reveal with shuffled presentation order.

4.2 Main Results

Table 1 presents the main experimental results comparing all four protocols. The first-arrival selection rate measures the probability that the customer agent selects the earliest-arriving proposal, with 33.3% representing unbiased (uniform) selection.

Under the HardGate baseline, the customer agent selects the first-arriving proposal 73.3% of the time—more than double the uniform expectation of 33.3%. This confirms severe first-proposal bias even when payment is blocked until all proposals are received. The sequential visibility of proposals through the message-fetching mechanism creates strong anchoring on the first-seen option.

Prompt-based interventions prove insufficient. SoftWait, which explicitly instructs the agent to wait for K proposals and compare before paying, achieves only a modest reduction to 63.3% first-arrival rate. Inference-Time Scaling, which samples $N = 5$ payment decisions and selects the lowest-price option, shows identical aggregate performance (63.3%). Neither approach eliminates the underlying bias mechanism.

EBR achieves a dramatic reduction to 24.4% first-arrival rate, representing a 48.9 percentage point improvement over HardGate ($z = 4.639$, $p < 0.001$). Critically, EBR's selection distribution is statistically indistinguishable from uniform (chi-squared test $p = 0.494$), indicating that arrival order no longer influences selection. Figure 2 visualizes the full rank distribution across protocols, showing that EBR achieves near-uniform selection (24.4%, 35.6%, 37.8% for ranks 1–3) while all baselines heavily favor the first-arriving proposal.

4.3 Per-Scenario Analysis

Table 2 breaks down first-arrival selection rates by scenario. HardGate shows extreme variance across scenarios (46.7% to 93.3%), while EBR maintains consistent near-uniform selection (13.3% to 33.3%) regardless of which contractor arrives first.

The ITS baseline exhibits particularly inconsistent behavior: it dramatically reduces bias in `contractors_first` (0.50 → 0.20) but worsens it in `contractors_third` (0.80 → 1.00).

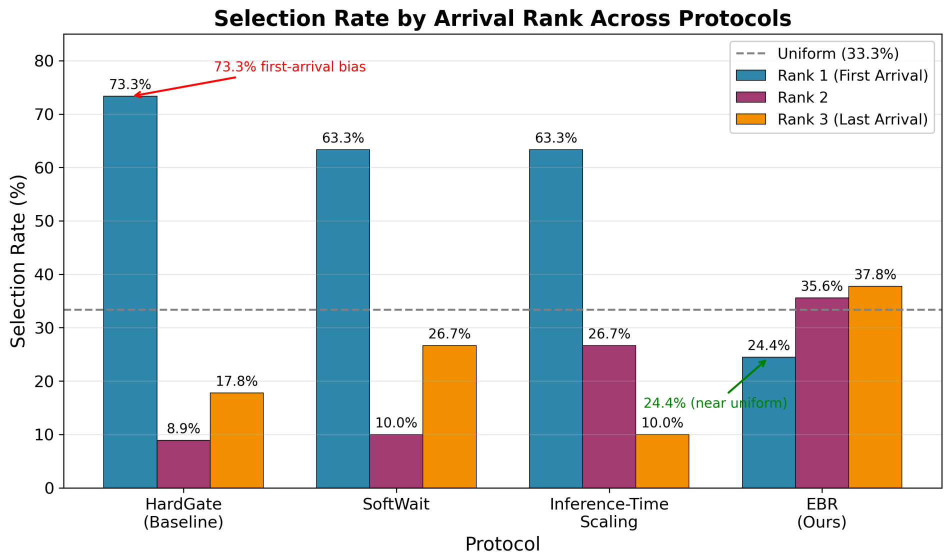

Figure 2: Selection rate by arrival rank across protocols. Under HardGate (baseline), the first-arriving proposal is selected 73.3% of the time, far exceeding the uniform expectation of 33.3% (dashed line). SoftWait and Inference-Time Scaling reduce bias slightly (63.3%) but remain far from uniform. EBR achieves near-uniform selection (24.4%, 35.6%, 37.8% for ranks 1–3), with chi-squared test confirming no significant deviation from uniform ($p = 0.494$).

Table 2: Per-scenario breakdown of first-arrival selection rates. Scenarios vary which contractor type arrives first. EBR shows consistent near-uniform selection across all scenarios, while baselines show high variance.

Protocol	contractors_first	contractors_second	contractors_third
HardGate	0.800	0.467	0.933
SoftWait	0.500	0.600	0.800
ITS	0.200	0.700	1.000
EBR (Ours)	**0.133**	**0.267**	**0.333**

This occurs because ITS selects the lowest-price option among proposals already visible at payment time, which may still favor early arrivers depending on the scenario's price structure. In contrast, EBR's consistent performance across scenarios demonstrates that its debiasing mechanism is robust to scenario-specific confounds.

4.4 Reveal-Position Analysis

While EBR eliminates arrival-order bias, we investigate whether bias simply shifts to presentation order within the revealed batch. Figure 3 shows selection rates by reveal position under EBR.

The first-revealed proposal (after shuffling) is selected 72.7% of the time (binomial test $p = 1.1 \times 10^{-7}$), virtually identical to HardGate's 73.3% first-arrival rate. This reveals that the LLM exhibits a fundamental primacy bias toward whatever proposal appears first in its context window, regardless of how that ordering was determined.

Critically, this does not diminish EBR's fairness value. Because EBR randomizes which proposal occupies the privileged first position, each business has an equal probability of being listed first over many transactions. EBR thus provides *statistical fairness* across the marketplace even though per-transaction primacy bias persists. This finding suggests that fully eliminating ordering bias in individual transactions may require additional interventions such as forced comparison steps or position-aware prompting.

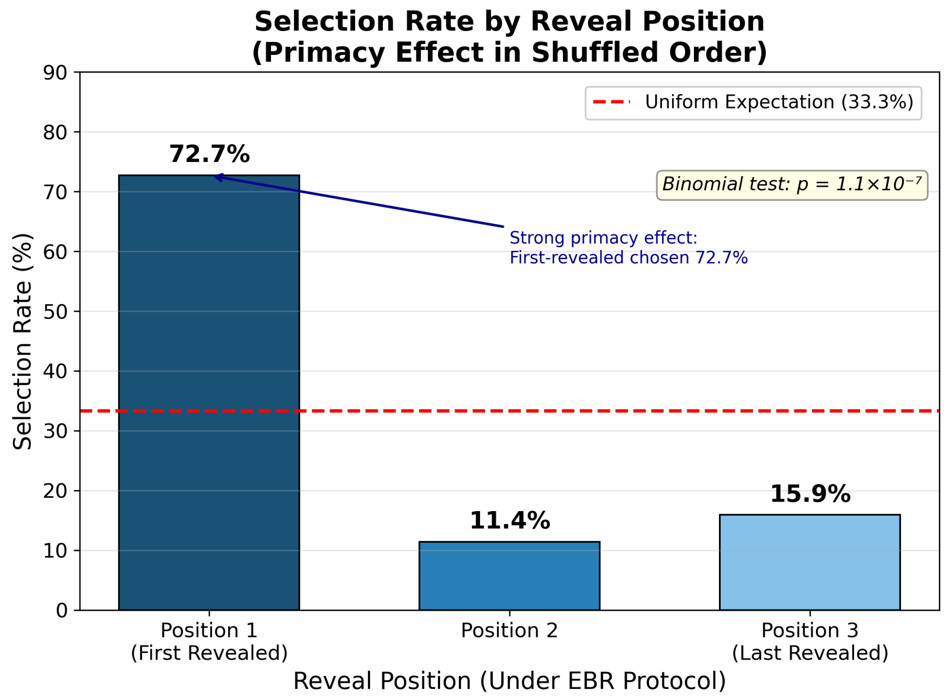

Figure 3: Selection rate by reveal position under EBR protocol. Despite shuffling arrival order, the first-revealed proposal is selected 72.7% of the time (binomial test $p = 1.1 \times 10^{-7}$), indicating a strong primacy effect in presentation order. This reveals that while EBR eliminates arrival-order bias, presentation-order bias persists—a finding with implications for future protocol design.

Table 3: Cross-model generalization of EBR effectiveness. EBR significantly reduces first-arrival bias on `gemini-2.5-flash` (48.9pp reduction, $p < 0.001$). On `claude-sonnet-4-5`, baseline bias is already near uniform (36.7%), creating a floor effect where further reduction is not statistically significant.

Model	HardGate	EBR	Δ	p-value
gemini-2.5-flash	0.733	0.244	**0.489**	<0.001***
claude-sonnet-4-5	0.367	0.233	0.133	0.13

4.5 Cross-Model Generalization

To assess generalization, we replicate the HardGate vs. EBR comparison on `claude-sonnet-4-5`. Table 3 presents the results.

On `claude-sonnet-4-5`, the HardGate baseline achieves 36.7% first-arrival rate—already near the uniform expectation of 33.3%. This model does not exhibit the satisficing behavior observed in `gemini-2.5-flash`; it naturally evaluates all available proposals before deciding. Consequently, EBR shows only a 13.3 percentage point reduction ($p = 0.13$), which is not statistically significant due to the floor effect.

This finding reveals that first-proposal bias is model-dependent. EBR is highly effective when the bias exists (as demonstrated on `gemini-2.5-flash`) but provides no additional benefit when the model already behaves rationally. The intervention should be recommended for models known to exhibit satisficing behavior, rather than as a universal protocol.

5 Conclusion

We introduced Escrowed Batch Reveal (EBR), a protocol-level intervention that eliminates first-proposal bias in LLM-mediated marketplaces. By buffering proposals and revealing them simultaneously in randomized order, EBR reduces first-arrival selection from 73.3% to 24.4% ($p < 0.001$), achieving statistical uniformity across arrival positions. Our results demonstrate that prompt-based interventions are insufficient—fairness requires architectural changes to information flow. While EBR successfully decouples selection from arrival timing, presentation-order primacy persists (72.7% first-revealed selection), indicating that complete debiasing remains an open challenge. These findings have immediate implications for the design of fair agentic economies.

References

Gagan Bansal, Wenyue Hua, Zezhou Huang, Adam Fourney, Amanda Swearngin, Will Epperson, Tyler Payne, Jake M. Hofman, Brendan Lucier, Chinmay Singh, Markus Mobius, Akshay Nambi, Archana Yadav, Kevin Gao, David Rothschild, Aleksandrs Slivkins, Daniel G. Goldstein, Hussein Mozannar, Nicole Immorlica, Maya Murad, Matthew Vogel, Subbarao Kambhampati, Eric Horvitz, and Saleema Amershi. Magentic marketplace: An open-source environment for studying agentic markets. *ArXiv*, abs/2510.25779, 2025.

Federico Bianchi, P. Chia, Mert Yüksekgönül, Jacopo Tagliabue, Daniel Jurafsky, and James Zou. How well can llms negotiate? negotiationarena platform and analysis. *ArXiv*, abs/2402.05863, 2024.

David Byrd, Maria Hybinette, and T. Balch. Abides: Towards high-fidelity market simulation for ai research. *ArXiv*, abs/1904.12066, 2019.

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*, abs/2310.05746, 2023.

J. Echterhoff, Yao Liu, Abeer Alessa, Julian McAuley, and Zexue He. Cognitive bias in decision-making with llms. pp. 12640–12653, 2024.

Shangmin Guo, Haoran Bu, Haochuan Wang, Yi Ren, Dianbo Sui, Yuming Shang, and S. Lu. Economics arena for large language models. *ArXiv*, abs/2401.01735, 2024.

Xiaobo Guo and Soroush Vosoughi. Serial position effects of large language models. *ArXiv*, abs/2406.15981, 2024.

Jiaxu Lou and Yifan Sun. Anchoring bias in large language models: an experimental study. *Journal of Computational Social Science*, 9, 2024.

J. Park, Joseph C. O'Brien, Carrie J. Cai, M. Morris, Percy Liang, and Michael S. Bernstein. *Generative Agents: Interactive Simulacra of Human Behavior*. 2023.

David Rothschild, Markus Mobius, Jake M. Hofman, E. Dillon, Daniel G. Goldstein, Nicole Immorlica, Sonia Jaffe, Brendan Lucier, Aleksandrs Slivkins, and Matthew Vogel. The agentic economy. *Communications of the ACM*, 69:39 – 42, 2025.

Eilam Shapira, Omer Madmon, Itamar Reinman, S. Amouyal, Roi Reichart, and Moshe Tennenholtz. Glee: A unified framework and benchmark for language-based economic environments. *ArXiv*, abs/2410.05254, 2024.

Nenad Tomasev, Matija Franklin, Joel Z. Leibo, Julian Jacobs, William A. Cunningham, Iason Gabriel, and Simon Osindero. Virtual agent economies. *ArXiv*, abs/2509.10147, 2025.

Felipe Valencia-Clavijo. Anchors in the machine: Behavioral and attributional evidence of anchoring bias in llms. *ArXiv*, abs/2511.05766, 2025.

H. Yin, Shai Vardi, and Vidyanand Choudhary. Fragile preferences: A deep dive into order effects in large language models. *ArXiv*, abs/2506.14092, 2025.

Stephan Zheng, Alexander R. Trott, Sunil Srinivasa, N. Naik, Melvin Gruesbeck, David C. Parkes, and R. Socher. The ai economist: Improving equality and productivity with ai-driven tax policies. *ArXiv*, abs/2004.13332, 2020.