

BID: BEHAVIORAL AGENTS IN DYNAMIC AUCTIONS

Weitong Zhang

Imperial College London
weitong.zhang@ic.ac.uk

Chengqi Zang

University of Tokyo
chengqi.zang@u-tokyo.ac.jp

Mark Schmidt

Naptha AI
mark@naptha.ai

Richard Blythman

Naptha AI
richard@naptha.ai

ABSTRACT

Complex societal systems are characterized by heterogeneous agents engaging in strategic interactions, yet current representative agent-based models (ABMs) struggle to capture these dynamics. We present BiD (Behavioral Agents in Dynamic Auctions), a novel ABM framework that focuses on modeling complex systems: heterogeneous agent modeling and socioeconomic dynamics. Using Dutch auctions as a microcosm, BiD models agent heterogeneity in risk preferences and dynamic trust scores, while modeling socioeconomic interactions via strategic communication phases before bidding. Our theoretical equilibria analysis reveals how BiD enables phenomena observed in real markets that are unexplained by classical ABMs, *i.e.*, successful low-valuation bidders win through strategic communication. We formalize behavioral agent strategies under different communication protocols and develop a reinforcement learning-guided policy for LLM-based agents to adapt their behaviors based on market dynamics. Experimental results demonstrate BiD’s capability in first modeling realistic market dynamics, providing socioeconomic perspectives for studying multi-agent systems and complex societal systems.

1 INTRODUCTION

Complex societal systems (*e.g.*, healthcare Boyd et al. (2022), social media Yang et al. (2024), economics Deguchi (2011), crypto Shannon et al. (2024), cities Batty (2007), and financial markets Cristelli (2013)) consist of many interconnected agents whose interactions create emergent behaviors. These behaviors cannot be predicted by analyzing individual actions alone Ladyman et al. (2013). Although conducting experiments with these systems is crucial in our digital age, real-world experiments are often costly or impossible An et al. (2021). Scientists thus rely on agent-based models (ABMs) to analyze real-world phenomena that empirical experiments cannot capture, particularly in understanding heterogeneous agent behaviors Zhang et al. (2021); Buchmann et al. (2016); Caiani et al. (2016) and socioeconomic dynamics Chen et al. (2023); Speybroeck et al. (2013); Axtell & Farmer (2022).

Representative ABMs, while effective in modeling basic agent interactions, struggle to capture the nuanced complexity of human behavior An et al. (2021). These models typically simplify agent behaviors to measurable thresholds, failing to address fundamental elements in context-dependent strategic decisions Wu et al. (2023), dynamic belief updating Zhang et al. (2024), and inter-agent communications Jin et al. (2024). The emergence of large language models (LLMs) has expanded the possibilities in this domain, demonstrating remarkable capabilities to mimic human-like behaviors Park et al. (2022); Zhou et al. (2023); Wang et al. (2023); Mou et al. (2024). These LLM-based agents can engage in sophisticated role-playing, participate in natural interactions with other agents Zhang et al. (2024), and execute complex decision-making tasks Park et al. (2022); Zhou et al. (2023) involving tool use Achiam et al. (2023). Researchers design social scenarios Yang et al. (2024) and distinct agent personas Jin et al. (2024), while developing scalable simulation platforms Yang et al. (2024) for real-world research and analysis Park et al. (2023). However, effectively integrating LLMs into ABMs requires moving beyond conventional benchmarks Ma et al. (2024). This integration is particularly urgent for modeling complex socioeconomic systems, where representative ABMs fall

short in capturing the interplay between heterogeneous agents and their strategic interactions under a realistic ABM Zhang et al. (2021); Ma et al. (2024).

Dutch auctions provide an ideal microcosm for studying complex strategic interactions for realistic ABMs. At system level, these descending price auctions exemplify key characteristics driving system inefficiencies: information asymmetry between buyers and sellers Maskin & Riley (2000); Guerci et al. (2014), allocation inefficiency from strategic waiting Gretschko & Wambach (2014), and agent behaviors that systematically deviate from theoretical predictions. Empirical evidence consistently shows that real-world outcomes diverge from oversimplified equilibrium predictions Bapna et al. (2004); Sow et al. (2012), *i.e.*, coordinated bidding in procurement auctions Patel (2021); Hortaçsu & Perrigne (2021), persistent price anomalies in agricultural markets Badau & Rada (2022). At agent level, homogeneity assumptions for agent characteristics have been challenged, including the agent’s prior heterogeneity Pham & Yamashita (2024) and distributional entry cost Moreno & Wooders (2011). Currently, existing tools face tractability constraints in modeling agent heterogeneity. Multi-agent simulations offer a promising alternative, capable of capturing complex patterns such as diverse risk preferences, trust scores, and agent-specific valuations. Thus, we propose BiD, a multi-agent-based ABM, to support research community studies in understanding Socioeconomic Dynamics, and Heterogeneous Agent Modeling.

Focus. Our work addresses the fundamental challenge of modeling complex socioeconomic systems through the lens of Dutch auctions. We propose BiD (Behavioral Agents in Dynamic Auctions), a novel ABM framework that advances conventional systems in two interconnected dimensions. First, we focus on capturing and analyzing socioeconomic dynamics in auction systems, particularly how information asymmetry, strategic communication, and trust relationships shape market outcomes. Using Dutch auctions as a microcosm, we demonstrate how these dynamics and interactions lead to empirically observed phenomena that deviate from classical predictions, such as successful low-valuation bidders and strategic waiting behaviors. Second, we develop a behavioral agent-based modeling approach that integrates heterogeneous agent characteristics with learning-based strategies. BiD combines heterogeneous agent modeling with reinforcement learning to capture how agents adapt their strategies based on heterogeneous characteristics, market dynamics, and historical interactions.

Contributions. (1) We demonstrate how BiD enables the exploration of previously unobservable decision-making processes in Dutch auctions, revealing strategic behaviors of low-valuation bidders that traditional economic models fail to capture. (2) By introducing a strategic communication mechanism, we provide a multi-agent framework for modeling heterogeneous agent interactions in complex socioeconomic systems. (3) Our approach bridges the gap between equilibrium predictions and empirical market dynamics, offering insights into how agents adapt their strategies under both private and public information.

Why Socioeconomic Dynamics matters and how BiD supports such analysis. BiD introduces strategic communication phases where agents exchange information strategically before bidding. This communication mechanism captures crucial socioeconomic dynamics missing from traditional models: in reality agents can signal, bluff, or cooperate to influence others’ decisions. Hautus et al. (2021) documented how communication among construction firms in Japanese procurement auctions led to systematic departures from efficient allocation, Robinson (1985) studied why certain auction mechanisms can lead bidders to form cartels to collude and some do not, but overall communications cannot be prevented. Our equilibrium analysis demonstrates how low-valuation bidders can win through strategic communication, a phenomenon observed in real markets but unexplained by classical theory. BiD’s reinforcement learning approach further allows agents to build strategies through experience, adapting their communication and bidding patterns based on market feedback and historical interactions.

How BiD works and why BiD supports Heterogeneous Agent Modeling. BiD formalizes a game of Multi-Round Mixed-Phase Dutch Auction (MRMP-DA), where agents compete through bidding and strategic communication phases. BiD models agent heterogeneity through risk preferences, trust scores, and belief updating mechanisms. These enable agents to exhibit diverse behaviors: risk-averse agents bid aggressively for assured wins, while risk-seeking agents strategically wait for higher expected surpluses. The trust scores update the belief based on the agents’ communication history and market dynamics, affecting how they perceive information from observation. Our results show that the heterogeneity leads to equilibrium behaviors that match empirical patterns in real markets.

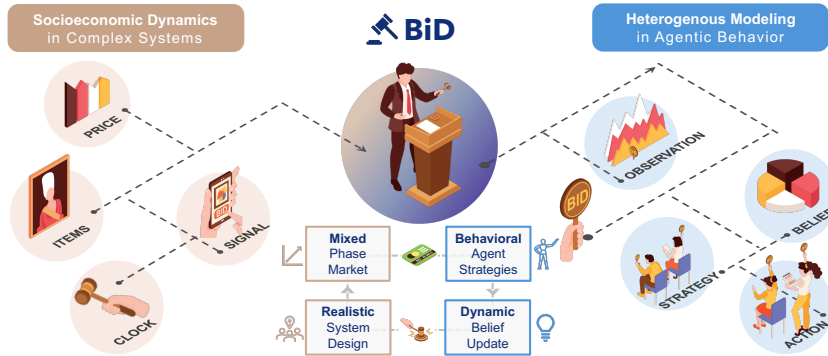


Figure 1: **BiD Illustration.** BiD bridges *socioeconomic dynamics* in complex systems and *heterogeneous modeling* in agentic behavior through strategic communication and diverse agent properties. As far as we know, this is the first time the Dutch auction has been studied in multi-agent systems and complex societal systems.

Insights. BiD’s heterogeneous agent modeling through risk preferences, trust scores, and action spaces under MRMP-DA, extends beyond auctions to broader market mechanisms including financial trading and resource allocation. The strategic communication phase and equilibrium analysis capture strategic interaction patterns in systems where agents can optimize decisions under incomplete information. Our dual focuses - heterogeneous agent modeling and socioeconomic dynamics - contribute to developing behavioral strategies in multi-agent systems, particularly for dynamic environments with strategic interactions.

2 SOCIOECONOMIC DYNAMICS OF DUTCH AUCTIONS

2.1 DUTCH AUCTION AS A GAME

A Dutch auction provides an ideal microcosm for studying complex strategic interactions ¹ In the case of selling, the auctioneer begins with a high asking price and lowers it with a fixed price reduction in each round until the participant accepts the price. We employ multi-unit ² Dutch auction where m rounds of the auction run for n agents, where $n - m = 1$, and the winning agent is eliminated from the bidders participating in the rest rounds; this setting is to encourage competition and prevent bidders from waiting until the last round to get the item for "free" and encourage more strategic interaction. e.g., low-valuation bidders can lie about their valuation to induce high-valuation bidder to accept the bids in earlier rounds, so that they have a better chance of winning in later rounds. When multiple bidders accept in a certain iteration, the tie-breaking rules will be in effect (e.g., random draw). Building on the multi-unit Dutch auction, we extend the classical setting by introducing strategic communication opportunities and trust dynamics to enable the study of more complex socioeconomic behaviors.

Game Formulation. A Multi-Round Mixed-Phase Dutch Auction (MRMP-DA) with heterogeneous agent can be formalized as a tuple $(\mathcal{N}, \Psi, \mathcal{S}, \Theta, \mathcal{H}, A, \mathcal{C}, U)$ where: \mathcal{N} represents the set of n agents participating in the m -round auction, $\Psi = \{\psi_{t,j} | j \in \{B, C\}\}$ defines the phases, where B represents bidding and C represents strategic communication in period t , where period refers to the cumulative number of phases, \mathcal{S} is the state space $\Theta = \times_{i=1}^n \Theta^i$ represents the joint type space, where Θ^i is agent i 's set of possible types (values) $A^i(\psi_{t,j})$ defines the available actions for agent i in phase $\psi_{t,j}$. \mathcal{H} represents the information states, where each agent has perfect recall of their previous states and actions $(h_0^i, a_0^i, h_1^i, a_1^i, \dots, h_t^i)$. \mathcal{C}_i^t is agent i 's vector of trust scores for other agents at period t , updated based on information state h_t^i and $u_i : \Theta \times \mathcal{S} \rightarrow \mathbb{R}^n$ maps agents' types and terminal states to utilities.

¹We use *Bidders* for socioeconomic dynamics to explore communication strategy. The next section details the transition from Bidders to Agents and BiD.

²multi-round where the auctions where items across rounds are homogeneous.

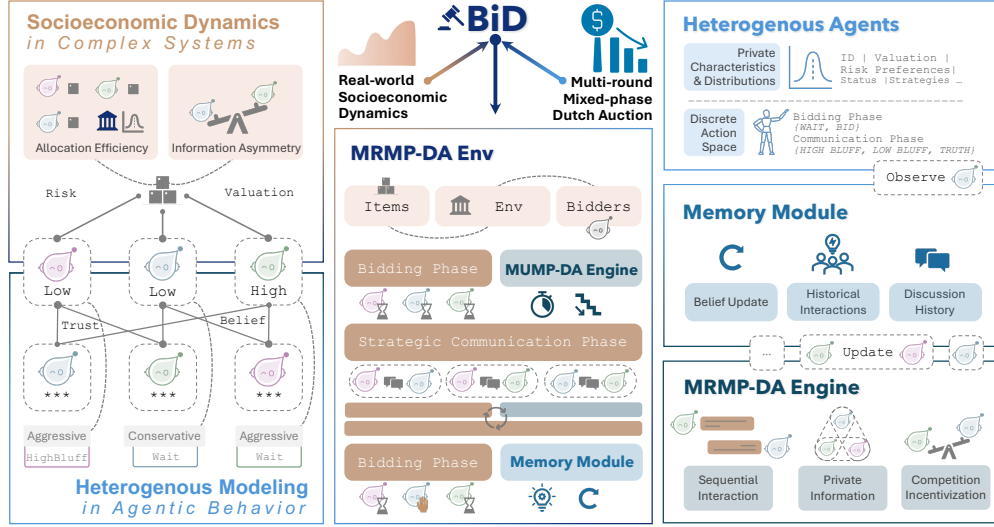


Figure 2: **BiD Design Overview.** A unified framework integrating *dynamic auction mechanism* (MRMP-DA Env), *strategic agent updates* (Memory Module), and *behavioral modeling* (Heterogeneous Agents) orchestrated by (MRMP-DA Engine) for *interaction and competition incentivization*. The RL-guided agents participate in MRMP-DA can align with real-world socioeconomic dynamics.

MRMP-DA Structure. MRMP-DA under imperfect-information Sevenster (2006); Friedman (2018) has three components Items, Bidders, and Environment, with two phase types, Bidding Phase and Strategic Communication Phase.

1. *Environment*. A rule-based env. that manages the auction flow, starting from price p_0 and applying fixed decrements until acceptance or reaching p_t .
2. *Bidders*. Bidders are heterogeneous and each of their type (valuation of item) is *private information*; their valuations are drawn from a continuous probability distribution $v_i \in f(v_i)$ that is *common knowledge* to all bidders. bidder i with valuation v_i in each round aims to maximize utility: $\max_{b_i} u_i(v_i - b_i, r_i)$ where b_i is the accepted price of bidder i , and u_i is a function that depends on the difference between valuation and we assume that each bidder's budget equals to her valuation. r_i is the risk-seekingness of bidder i , higher r_i signifies that i is more risk-loving, we use the functional form $u_i = e^{r(v_i)(v_i - b_i)}$ throughout such that $r(\cdot)$ is a weakly increasing function of v_i .
3. *Items*. Items for auction are homogeneous in our multi-around descending auction, the homogeneity indicates that in each round, the bidder's private valuation toward the item does not change.
 1. *Bidding Phase*. Each round for the bidding phase, the auctioneer start from a high selling price p_0 , then, for each iteration, the auctioneer call a fixed decrement to the price of the item, such that $p_t - p_{t-1} = p_{t-1} - p_{t-2}$ for all discrete $t \in [0, T]$, the current round ends when some bidder accepts the price p_t , the auction be deemed unsuccessful if no one accepts the price until p_T , however this does not happen in our setting as accepting the price p_T yields a positive payoff for any bidder as we assumed that $p_T < f(v_i) < p_0$, such that f is the probability distribution of agents' valuation.
 2. *Strategic Communication Phase*. We also introduce a strategic communication phase. Bidders are randomly matched in pairs before or during bidding phase, and when that happens, matched agents can (simultaneously or sequentially) freely choose to reveal any information (truth or lie) with the matched bidder with the goal of maximizing utility calculated by u_i . After the strategic communication phase is over, the game then is proceeded back to the bidding phase such that the auctioneer continue to run price decrement following from where the auction was left off.

Trust Score. Among bidders, each one keeps track of trust scores for other bidders in the form of a vector, e.g. bidder i 's credit score vector is given by $\mathcal{C}^i = \{c_{i,1}, c_{i,2}, \dots, c_{i,i-1}, c_{i,i+1}, \dots, c_{i,n}\}$, the trust scores bidder i assigns to other bidders. $c_{i,j}$ is affected when the bidder j does/does not obey

the information or proposed action she conveyed during the strategic communication phase. This is treated as a variant of belief updates across episodes.

While a complete equilibrium analysis of a full-fledged BiD model would be intractable, understanding its theoretical advantages is crucial. By examining the equilibrium properties of the final round, we demonstrate how communication-enabled heterogeneous agents achieve superior allocation efficiency and welfare outcomes compared to traditional representative agent-based models.

2.2 REPRESENTATIVE ABM v.s. BiD

A Single-Round One-Phase Representative ABM Model. We analyze a no-strategic communication, representative-agent auction that serve as the benchmark system. Under our formulation, such an traditional Dutch auction can be seen as a Single-Round One-Phase (Bidding only) auction for two bidders,

Proposition 2.1 *A pure strategy Nash equilibrium strategy for the auction above is given by $b^*(v_i, r_i) = v_i - 1/r_i$.*

Since r_i (risk-seekingness) weakly increase with v_i , we have that b^* is strictly increasing in v_i , which means that **low-valuation bidder never wins**, however since dutch auctions run fixed decrement and discrete biddable prices, low valuation bidders win through tie-breaking mechanisms if we constrain strategies to be pure. However, empirical evidence suggests that outcomes frequently deviate from this theoretical prediction in real scenarios and bidders’ behaviors can hardly be regulated to be ‘no communication’. These observations motivate us to introduce a communication mechanism that captures this prevalent feature of real-world auctions. Specifically, our mixed-phase framework where bidders can engage in private communication before submitting their bids, allowing strategic information sharing and signaling that may influence final allocations. This extension not only better aligns our theoretical model with observed market behavior but also provides a rich setting for analyzing how information exchange affects outcomes, efficiency, and strategic interactions.

A Single-Round Two-Phase BiD Model. To draw comparison, we analyze the last round of a MRMP-DA with two bidders, which is a Single-Round Two-Phase Dutch Auction. We name them bidder 1 and bidder 2, without loss of generality; we assume $v_1 > v_2$, and before communication, the two bidders’ common knowledge is that their valuations are drawn from a uniform distribution $U[0, 1]$; during strategic communication, bidder 1 reports her valuation first, and bidder 2 reports conditioned on bidder 1’s report. After strategic communication ends, the Dutch auction starts and ends when one bidder accepts the bid. Before diving into the main existence theorem, we provide a proposition when bidders are at the strategic communication phase.

Proposition 2.2 *Over-reporting in a single-round Dutch auction is a weakly dominated strategy for any bidder.*

Proof: Since the bidder’s valuation is drawn from a continuous distribution, we have $\mathbf{P}(v_i = v_j) = 0$. Therefore, we consider two cases of bidder j over-reporting to bidder i , when $v_i > v_j$ and $v_i < v_j$, about her valuation $v_{j,i}$, the notation denotes that direction of message $j \rightarrow i$. We will mainly analyze how over-reporting would change opponents’ behavior.

Case 1. $v_i > v_j$, if bidder j over-report his valuation to i such that we denote as $v_{j,i} > v_j$, then for bidder i to secure a higher payoff, bidder i accept the price at $p_t > v_{j,i}$ with higher probability since he is afraid that bidder j might steal the item, which as a result bidder j has less possibility of winning the item.

Case 2. $v_i < v_j$, even if bidder j over-report her valuation, bidder i would not accept any $p_t > v_i$ since that would yield negative payoff, therefore bidder i would not change her bidding strategy.

Thus proposition gives the equilibrium belief update in the upcoming existence theorem of a Perfect Bayesian Equilibrium(PBE) in our Single-Round Two-Phase game.

Theorem 2.3 *(Existence of Perfect Bayesian Equilibrium.) A PBE exists in the Single-Round Two-Phase BiD model where the equilibrium belief update is given by $0 < c_i < 1$ that the opponent is telling the truth and $1 - c_i$ that the opponent is lying based on Prop. 2.2; and c_i here correspond to the trust score we mentioned in Sec. 2. Furthermore, this equilibrium utility is **weakly higher** than the equilibrium utility in Single-Round One-Phase case.*

Thm. 2.3 ensures the existence of an equilibrium under a sophisticated and realistic setting. In Table. 1 we numerically solve for this equilibrium and validated that this equilibrium elicit strategic interaction and welfare gain; and secondly, with a full-fledged BiD model in Sec. 3., we are able to simulate cases that matches realistic auction markets in Fig. 3.

3 MODELING AND LEARNING HETEROGENEOUS AGENT STRATEGIES

3.1 WORKFLOW OF BID

MRMP-DA Env under Socioeconomic Dynamics. It defines the fundamental structural parameters, including the number of agents, items, and auction rounds following MRMP-DA Structure in Sec. 2.1 for socioeconomic dynamics. The environments’ functionalities include tracking auction state progression, managing agent participation status, and calculating reward structures. The environment generates initial auction conditions, determines agent elimination rules, and maintains an agent state representation that supports market dynamics.

Heterogeneous Agents with Behavioral Strategies. We design a computational representation of heterogeneous auction participants. Each agent is characterized by unique and private valuation distributions, risk preferences, and adaptive behavioral strategies. Agents operate within a discrete action space comprising five fundamental actions: WAIT, BID, HIGH_BLUFF, LOW_BLUFF, and TRUTH, governed by a behavioral strategy that integrates belief modeling and strategic decision-making.

Agent i ’s behavioral strategy in phase $\psi \in \Psi$ is defined as a mixed strategy $\pi^i(\psi_{t,j}) \in \Delta(A^i(\psi_{t,j}))$ which is a probability distribution over available actions given a phase. The reason we call it ‘behavioral’ is that we programmed the agents to deviate from the strictly rational strategy with some randomness. And $\pi_{\psi}^i(a^i|h^i)$ denotes the probability that agent i takes action a^i in phase ψ given information state h^i . Agent i ’s strategy consists of belief modeling and action selection and is expressed as

$$\pi_{\psi}^i(a^i|h^i) = \sum_{\theta \in \Theta} b^i(\theta|h^i, \mathcal{C}^i) \tilde{\pi}_{\psi}^i(a^i|h^i, \theta, \mathcal{C}^i) \quad (1)$$

where $\tilde{\pi}_{\psi}^i$ is a belief-conditioned behavioral strategy for agent i in phase ψ . Denote $\pi = (\pi^1, \dots, \pi^n)$ as a collection of strategies of all agents, and the expected utility for agent i induced by strategy profile π as $\mathbb{E}_{\pi}[u_i]$.

Memory Module for Belief Update. Implemented through a Bayesian learning framework Dekel et al. (2004); Huang & Zhu (2019); Zhang et al. (2024), the memory enables agents to construct and update beliefs about other agents’ types. This module maintains communication history, action records, and dynamic trust scores, allowing for tractable belief updates based on historical interactions.

Given the imperfect information nature of MRMP-DA, agent i form a belief $b^i : H^i \rightarrow \Delta(\Theta)$ on all agents’ types based on its observation. We define agent i ’s belief on information state $h^i \in H^i$ as:

$$b^i(\theta|h^i) \stackrel{def}{=} \frac{p^i(\theta)p^i(h^i|\theta)}{\sum_{\theta' \in \Theta} p^i(\theta')p^i(h^i|\theta')} \quad (2)$$

where $p^i(\theta)$ is the prior probability of all agents’ types from agent i ’s view, and $p^i(h^i|\theta)$ is the probability that agent i observes h^i given joint types θ . For each agent i , the trust score for agent j at time $t + 1$ can be expressed as:

$$\mathcal{C}_{t+1}^i[j] = f(\mathcal{C}_t^i[j], h_t^i) \quad (3)$$

where f is a trust update that considers the current trust score and observed history. Let $\Delta h_t^i = h_{t+1}^i \setminus h_t^i$ denote the new information contained in h_{t+1}^i at step $t + 1$ compared to h_t^i , then agent i ’s belief can be updated via Bayes’ rule:

$$b_{t+1}^i(\theta|h_{t+1}^i, \mathcal{C}_{t+1}^i) = \frac{b_t^i(\theta|h_t^i, \mathcal{C}_t^i)p^i(\Delta h_t^i|\theta)}{\sum_{\theta' \in \Theta} b_t^i(\theta'|h_t^i, \mathcal{C}_t^i)p^i(\Delta h_t^i|\theta')} \quad (4)$$

And $b_0^i(\theta|h_0^i) = b_0^i(\theta)$ is set as a uniform distribution, which is an unbiased estimate if there is no prior information.

MRMP-DA Engine for *Efficiency Incentives*. orchestrates the strategic interaction mechanism, managing the descending price auction dynamics, and facilitating the strategic communication phases. The engine designs communication protocols that incentivize strategic signaling while maintaining private information. The engine compels agents to deviate from strictly rational decision-making patterns.

MRMP-DA Engine orchestrates the interaction mechanism in a sequential manner: without informing either agent, the agent with higher valuation is always selected to report first. This design engages dynamic in ABM where the second agent can condition their strategy on the first agent’s reported valuation. Thus, agents incorporate different strategic choices within a discrete action space Z . The Engine extend the standard behavioral strategy by conditioning it on both the agent’s beliefs and their strategic choices:

$$\pi_\psi^i(a|h) = \sum_{\theta \in \Theta} b^i(\theta|h) \sum_{z \in Z} \mu^i(z|h, \theta) \tilde{\pi}_\psi^i(a|h, \theta, z) \quad (5)$$

$b^i(\theta|h)$ represents the agent’s belief function, $\mu^i(z|h, \theta)$ denotes the strategy selection policy, and $\tilde{\pi}_\psi^i(a|h, \theta, z)$ is the belief-conditioned behavioral strategy. This formulation engages agents to select actions based on their private information and beliefs about other Heterogeneous Agents’ valuations in MRMP-DA Env, either participate directly in the auction (WAIT, BID) or engage in strategic sequential communication (HIGH_BLUFF, LOW_BLUFF, TRUTH) that influences other agents’ beliefs and subsequent actions.

3.2 RL-GUIDED STRATEGY LEARNING

Building upon the theoretical foundations of BiD and the MRMP-DA framework, we implement a RL approach that leverages Conservative Q-Learning (CQL) Kumar et al. (2020) to learn optimal bidding and communication strategies. Our implementation addresses three key challenges: (1) the complex mixed-phase nature of MRMP-DA, (2) the need to maintain strategic diversity while optimizing behavior, and (3) the requirement for robust offline learning from demonstration data. We first generate auction trajectories using LLM-guided actions, creating a dataset $\mathcal{D} = \{s, z, r, s'\}$ that captures the full range of strategic interactions. Different from Kumar et al. (2020); Jin et al. (2024), the state space s encompasses both explicit auction parameters (e.g., current price, remaining items) and implicit strategic information (e.g., trust scores, belief distributions). The CQL loss function is formulated as:

$$\begin{aligned} \mathcal{L}(\phi) = & \rho \mathbb{E}_{s \sim \mathcal{D}} \left[\log \sum_z \exp Q_\phi(s, z) - \mathbb{E}_{z \sim \mathcal{D}} [Q_\phi(s, z)] \right] \\ & + \frac{1}{2} \mathbb{E}_{s, z, s' \sim \mathcal{D}} [(Q_\phi - \mathcal{B}Q_\phi)^2] + \lambda \mathcal{R}(\phi) \end{aligned} \quad (6)$$

where ρ controls the conservatism-optimality trade-off, \mathcal{B} represents the Bellman operator, and we introduce a regularization term $\mathcal{R}(\phi)$ to maintain strategic diversity. The hyperparameter λ balances the importance of strategic diversity against pure reward maximization. The learned policy enables agents to execute sophisticated strategic behaviors that align with our theoretical equilibrium analysis while adapting to the dynamic nature of practical auction settings. To ensure robust learning, we implement both episodic and step-based evaluation metrics, tracking not only traditional RL metrics (TD error, cumulative rewards) but also auction-specific indicators, *i.e.*, price efficiency and trust evolution.

4 EXPERIMENT

We conduct three experiments to evaluate the performance of our proposed BiD, as a strategic agent-based system for analyzing dynamic auction: (1) numerically computing and analyzing strategic communication’s impact on bidders’ equilibrium utility in a two-bidder setting of Thm. 2.3, (2)

Table 1: **Convergence Utility Comparison.** Rep. ABM without (W/O) Communication vs. BiD with (W/) Strategic Communication between High (H-) Low (L-) value bidders on utility (Util.) convergence.

High Value	Low Value	Rep. ABM: W/O Communication			BiD: W/ Communication		
Value	Value	H-Util.	L-Util.	L-Win Rate	H-Util.	L-Util.	L-Win Rate
0.50	0.10	0.9363	0.0000	0.00%	1.0242	0.0722	7.22%
0.54	0.14	1.0494	0.0002	0.02%	1.1487	0.1146	9.85%
0.59	0.19	1.1883	0.0002	0.02%	1.3184	0.1171	10.65%
0.63	0.23	1.3768	0.0090	0.89%	1.5470	0.1644	11.81%
0.68	0.28	1.6302	0.1565	15.28%	1.8514	0.2146	13.05%
0.72	0.32	1.9711	0.2142	20.69%	2.2582	0.2667	14.32%




Figure 3: **Bidding Behavior Distribution.** Empirical data (778 NFT Dutch auction bids, over 8 weeks) are captured by BiD’s heterogeneous agents with regard to diverse human behaviors and transaction distributions. Real-market socioeconomic dynamics are reproduced through 50 multi-round mixed-phase Dutch auctions. (Sharp similarities, zoom in for more details.)

Table 2: **Strategic Analysis Across Agents.** Quantitative evaluation of auction performance through (1) observation conditions {W/, W/O prior}, (2) communication phases {W/, W/O strategic communication}, and (3) agent frameworks {Representative (Rep.) ABM, Heterogeneous (Heterg.) ABM, BiD}, measured by reward distribution, bidding dynamics, and market efficiency.

Agents	Metrics (Avg.)	Without Communication		With Communication	
		Rep. ABM (W/O Prior)	Rep. ABM (W/ Prior)	Heterg. ABM LLM-guided	BiD (Ours)
GPT-4o	Total Rewards	185.88	202.90	435.25	450.28
	Bid round	4.4	4.5	6.3	6.5
	H-Rewards	24.31	25.61	51.98	54.11
	L-Rewards	6.54	14.90	16.33	18.24
4o-Mini	Total Rewards	172.31	186.21	301.00	293.24
	Bid round	4.6	4.9	5.1	5.2
	H-Rewards	25.67	28.83	35.82	38.59
	L-Rewards	4.53	4.69	16.18	17.24
3-Sonnet	Total Rewards	89.72	133.74	186.66	201.09
	Bid round	4.6	4.6	4.8	5.6
	H-Rewards	16.12	28.82	25.17	33.10
	L-Rewards	Fail	8.30	4.53	11.80
3-Opus	Total Rewards	174.56	193.29	233.47	282.67
	Bid round	5.4	4.1	5.5	6.4
	H-Rewards	24.45	26.09	43.02	53.70
	L-Rewards	8.02	7.96	13.8	18.02

demonstrating its generalizability by studying real-world NFT Dutch auction data which reveals complex socioeconomic phenomena, and (3) evaluating the system’s scalability with heterogeneous agents, particularly focusing on how low-valuation bidders can compete through strategic communication.

Setup. We conduct experiments using OpenAI {GPT-4o, 4o-mini} Achiam et al. (2023), Claude {3-Sonnet, 3-Opus} Anthropic (2024) for generating strategic communications under the MRMP-DA structure. For modeling and learning the bidding and communication strategy, we collect interaction trajectories as our offline dataset. We implement improved CQL for policy learning (learning_rate:5e-5, Step: 50000, Critic: 2, Interval: 1000, batch_size: 64, Alpha: 2.5). We evaluate several variants of bidding agents: (1) BiD: our full-fledged model with RL-guided strategy and LLM-based framework with MRMP-DA setting, (2) LLM-guided: initialized agents using LLM for both strategy selection and communication, can be seen as a simplified Heterogeneous ABM, (3) Representative ABM W/ W/O prior: benchmark agents W/ WO common prior, while they can not conduct strategic communication as Sec. 2.2. Each experiment is repeated 25 auctions for statistical significance.

Equilibrium Analysis. We first examine the impact of strategic communication in a two-bidder setting in Table 1. This allows us to observe how strategic communication mechanism affects bidders’ equilibrium utility and also validates the existence of equilibria where LLM-based agents can approximate and converge to.

In the benchmarking representative ABM without communication, low-value bidders consistently achieve near-zero utility (0.0000-0.2142) even if we allow for a mixed strategy, while high-value bidders maintain high utilities (0.9363-1.9711). This aligns with our theoretical prediction in Prop. 2.1 where low-value bidders are disadvantaged. However, with BiD’s strategic communication mechanism, we observe a consistent improvement in low-value bidder utilities (0.0722-0.2667), and even better outcomes for high-value bidders (1.0242-2.2582). This equilibrium convergence supports Thm. 2.3’s prediction that strategic communication can lead to higher equilibrium utilities.

The last two rows of Table. 1 in low valuation bidder’s win rate may seem disadvantageous for the ‘w/ communication’ group, however, they actually exemplifies how strategic interaction and risk-seeking bidding behaviors leads to higher surplus ($v_i - b_i$), thus resulting in higher utility with lower win rate; on the other hand, we can see that the ‘w/o communication’ low valuation bidder hardly extract

positive surplus from winning the bid because $\frac{utility}{win-rate} \rightarrow 1$, which means that they win mostly by bidding their own valuations.

Market Dynamics Analysis. To validate BiD’s generalizability and plausibility in capturing real-world market dynamics, we analyze a global NFT (non-fungible token) Dutch auction dataset Shannon et al. (2024). The transparent yet pseudonymous nature provides a unique window for studying market interactions from real strategic wallet behaviors while maintaining privacy. The ‘World of V’ NFT dataset comprises 778 Dutch auction sales over an 8-week period, with sale prices ranging from 537.44 VET (price) to 49,513.65 VET. We segmented the sales data into three subgroups by the final sales price in VET to equal number of samples and plotted the Proportion of Valuation Retained, which is given by $\frac{b-p_T}{p_0-p_T}$, such that b is the successful bid price, p_0, p_T are start and end price of the auction.

Surprisingly, the three subgroups exhibit very distinct bidding behaviors. For the high price segment, successful prices are concentrated at high retention rate, which reflected the strong consensus in valuation for expensive items and the fear of losing it, conversely, for the low price segment, bids cluster at 20% - 40% retention rate, which shows that bidders are more willing to wait for a better deal rather than winning the item; whereas for the middle price segment, they exhibit properties of both H- and L- price segment and shows high frequencies on both ends. With full-fledged BiD model, we are able to simulate the distribution that matches the real distribution for all price segments.

Strategic Behavior Analysis. Having validated BiD’s capability in modeling real-world market dynamics, we now evaluate its generalization capabilities and robust performance in capturing complex auction dynamics (Table 2).

Notably, BiD outperforms both rep. ABMs and heterg. ABM benchmarks across all metrics and configurations. The higher total rewards achieved by BiD (450.28, GPT-4o) compared to rep. ABMs (202.90, GPT-4o) shows the theoretical predictions from Thm. 2.3 regarding the benefits of strategic communication. The increased number of bidding rounds (6.5 for BiD v.s. 6.3, 4.5, 4.4 for benchmark models) also supports our market dynamics analysis in Fig. 3, suggesting more strategic waiting behaviors and complex bidding patterns. Moreover, the dual improved distribution of rewards for both high-value (38.59, 4o-Mini) and low-value (17.24, 4o-Mini) agents under BiD, compared to the stark disparity in rep. ABMs (H: 28.83, L: 4.69), empirically supports our equilibrium analysis from Table 1. L-agents have failed to bid on some occasions. This demonstrates how BiD enables low-value bidders to achieve better outcomes without compromising the utilities of high-value bidders, ultimately contributing to improved allocation efficiency and global rewards.

5 DISCUSSION

Socioeconomic Dynamics in Practice. BiD’s architecture emerge from real-world strategic interactions observed in various market mechanisms. Our framework captures essential behavioral patterns through the integration of strategic communication phases, trust scoring, and belief updates - elements that naturally arise in empirical market settings. The modular design allows BiD to model complex phenomena like information cascades, strategic misrepresentation, and trust network formation that are prevalent in real-world auctions but often overlooked in classical frameworks. This grounding in empirical observations enables BiD to replicate and analyze market behaviors that emerge from agent interactions rather than being explicitly programmed.

Scalable Heterogeneous Agent Modeling. BiD achieves scalability by focusing on essential characteristics of Dutch auctions, while preserving key strategic elements. The framework abstracts away non-critical auction parameters such as item heterogeneity, temporal dynamics, bidding increment variations, and entry/exit mechanisms, while maintaining crucial features like imperfect information and private valuations. This selective abstraction, combined with our modular agent architecture, enables BiD to efficiently scale to large numbers of agents without compromising the fundamental strategic dynamics. The framework’s ability to handle heterogeneous risk preferences, trust relationships, and belief systems demonstrates its capability for modeling diverse agent populations in complex market environments.

RL-guided Agent Strategy Emergence. The learned policies exhibit four key strategic adaptations that naturally emerge during learning: (1) value-conditional communication patterns, where agents’

communication strategies adapt based on their position in the value distribution, (2) dynamic bid timing in response to competitive pressure and trust network evolution, (3) trust-building behaviors where agents sacrifice immediate gains for future cooperation, and (4) context-aware risk adjustment based on auction progression. Our RL-guided agents autonomously discover many of the equilibrium properties. BiD’s formulation is able to capture essential strategic dynamics.

Strategic Communication as an Engine. Our research demonstrates that strategic communication is not a peripheral feature, but a fundamental dynamic that reshapes agent interactions and outcomes of an agent network. By modeling communication protocols, we reveal how heterogeneous agents can strategically exchange information, dynamically update trust, and ultimately alter allocation efficiency. This approach transcends traditional equilibrium and representative models, showing that designed communication is a strategic instrument through which agents can adaptively negotiate, signal, and optimize their outcomes in complex socioeconomic systems. We further explore diverse BiD variants of strategic communication in Appendix A.2.

Limitation. While BiD models strategic behaviors in Dutch auctions, the framework’s reliance on pairwise communication for equilibrium analysis may not fully capture the complex network effects present in large-scale markets where information diffuses through multiple channels simultaneously. This abstraction, though enabling tractable analysis, points to opportunities for extending the framework to model large-scale propagation mechanisms in future work.

6 CONCLUSION

BiD is an agent-based modeling framework designed to capture complex strategic interactions in Dutch auctions through heterogeneous agent modeling and socioeconomic dynamics, it demonstrates how integrating strategic communication with behavioral modeling can better explain empirically observed market phenomena, such as low-valuation winners and dynamic trust-based strategies. Our theoretical analysis and experimental results validate BiD’s ability to model realistic market dynamics while maintaining computational tractability. The framework’s success in capturing emergent behaviors suggests its potential applications beyond auctions to broader market mechanisms where strategic communication and trust dynamics are crucial.

IMPACT STATEMENT

This paper presents research on modeling complex strategic agent interactions in dynamic auctions. While our work advances understanding of multi-agent systems and market dynamics, we acknowledge that improved modeling of strategic behaviors could have broader societal implications. However, as our framework focuses on analyzing and understanding existing agentic behaviors rather than manipulating financial market outcomes, we believe the primary impact will be academic and analytical in nature.

ACKNOWLEDGMENT

This work was supported by the JADS programme and UK Research and Innovation [UKRI Centre for Doctoral Training in AI for Healthcare grant number EP/S023283/1].

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Li An, Volker Grimm, Abigail Sullivan, BL Turner Ii, Nicolas Malleson, Alison Heppenstall, Christian Vincenot, Derek Robinson, Xinyue Ye, Jianguo Liu, et al. Challenges, tasks, and opportunities in modeling agent-based complex systems. *Ecological Modelling*, 457:109685, 2021.
- Anthropic. The claude 3 model family: Opus, sonnet, haiku. *Online*, <https://www.anthropic.com>, 2024.
- Robert L Axtell and J Doyne Farmer. Agent-based modeling in economics and finance: Past, present, and future. *Journal of Economic Literature*, pp. 1–101, 2022.
- Flavius Badau and Nicholas Rada. Disequilibrium effects from misallocated markets: An application to agriculture. *Agricultural Economics*, 53(4):592–604, 2022.
- Ravi Bapna, Paulo Goes, Alok Gupta, and Yiwei Jin. User heterogeneity and its impact on electronic auction market design: An empirical exploration. *MIS quarterly*, pp. 21–43, 2004.
- Michael Batty. *Cities and complexity: understanding cities with cellular automata, agent-based models, and fractals*. The MIT press, 2007.
- Jennifer Boyd, Rebekah Wilson, Corinna Elsenbroich, Alison Heppenstall, and Petra Meier. Agent-based modelling of health inequalities following the complexity turn in public health: a systematic review. *International Journal of Environmental Research and Public Health*, 19(24):16807, 2022.
- Carsten M Buchmann, Katrin Grossmann, and Nina Schwarz. How agent heterogeneity, model structure and input data determine the performance of an empirical abm—a real-world case study on residential mobility. *Environmental modelling & software*, 75:77–93, 2016.
- Alessandro Caiani, Alberto Russo, Antonio Palestrini, and Mauro Gallegati. Economics with heterogeneous interacting agents. *New Economic Windows, Springer Series. DOI*, 10:978–3, 2016.
- Yufeng Chen, Liyan Xu, Xiao Zhang, Zilin Wang, Hailong Li, Yansheng Yang, Hong You, and Dihua Li. Socio-econ-ecosystem multipurpose simulator (seems): An easy-to-apply agent-based model for simulating small-scale coupled human and nature systems in biological conservation hotspots. *Ecological Modelling*, 476:110232, 2023.
- Matthieu Cristelli. *Complexity in financial markets: modeling psychological behavior in agent-based models and order book models*. Springer Science & Business Media, 2013.
- Hiroshi Deguchi. *Economics as an agent-based complex system: toward agent-based social systems sciences*. Springer Science & Business Media, 2011.

- Eddie Dekel, Drew Fudenberg, and David K Levine. Learning to play bayesian games. *Games and economic behavior*, 46(2):282–303, 2004.
- Daniel Friedman. The double auction market institution: A survey. In *The double auction market*, pp. 3–26. Routledge, 2018.
- Irving L Glicksberg. A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.
- Vitali Gretschko and Achim Wambach. Information acquisition during a descending auction. *Economic Theory*, 55:731–751, 2014.
- Eric Guerci, Alan Kirman, and Sonia Moulet. Learning to bid in sequential dutch auctions. *Journal of Economic Dynamics and Control*, 48:374–393, 2014.
- Michael J Hautus, Neil A Macmillan, and C Douglas Creelman. *Detection theory: A user’s guide*. Routledge, 2021.
- Ali Hortaçsu and Isabelle Perrigne. Empirical perspectives on auctions. In *Handbook of Industrial Organization*, volume 5, pp. 81–175. Elsevier, 2021.
- Linan Huang and Quanyan Zhu. Dynamic bayesian games for adversarial and defensive cyber deception. *Autonomous Cyber Deception: Reasoning, Adaptive Planning, and Evaluation of HoneyThings*, pp. 75–97, 2019.
- Xuanfa Jin, Ziyang Wang, Yali Du, Meng Fang, Haifeng Zhang, and Jun Wang. Learning to discuss strategically: A case study on one night ultimate werewolf. *arXiv preprint arXiv:2405.19946*, 2024.
- Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 33:1179–1191, 2020.
- James Ladyman, James Lambert, and Karoline Wiesner. What is a complex system? *European Journal for Philosophy of Science*, 3:33–67, 2013.
- Qun Ma, Xiao Xue, Deyu Zhou, Xiangning Yu, Donghua Liu, Xuwen Zhang, Zihan Zhao, Yifan Shen, Peilin Ji, Juanjuan Li, et al. Computational experiments meet large language model based agents: A survey and perspective. *arXiv preprint arXiv:2402.00262*, 2024.
- Eric Maskin and John Riley. Asymmetric auctions. *The review of economic studies*, 67(3):413–438, 2000.
- Diego Moreno and John Wooders. Auctions with heterogeneous entry costs. *The RAND Journal of Economics*, 42(2):313–336, 2011.
- Xinyi Mou, Zhongyu Wei, and Xuanjing Huang. Unveiling the truth and facilitating change: Towards agent-based large-scale social movement simulation. *arXiv preprint arXiv:2402.16333*, 2024.
- Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Social simulacra: Creating populated prototypes for social computing systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, pp. 1–18, 2022.
- Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pp. 1–22, 2023.
- P Patel. Modelling cooperation, competition, and equilibrium: The enduring relevance of game theory in shaping economic realities. *Social Science Chronicle*, 1:1–19, 2021.
- Hien Pham and Takuro Yamashita. Auction design with heterogeneous priors. *Games and Economic Behavior*, 145:413–425, 2024.
- Marc S Robinson. Collusion and the choice of auction. *The RAND Journal of Economics*, pp. 141–145, 1985.

- Halsey Royden and Patrick Michael Fitzpatrick. *Real analysis*. China Machine Press, 2010.
- Merlijn Sevenster. *Branches of imperfect information: logic, games, and computation*. University of Amsterdam, 2006.
- Darren Shannon, Michael Dowling, Marjan Zhaf, and Barry Sheehan. Dutch auction dynamics in non-fungible token (nft) markets. *Economic Modelling*, 141:106882, 2024.
- Jacob Sow, Patricia Anthony, and Chong Mun Ho. Homogeneous and heterogeneous agents in electronic auctions. In *Practical Applications of Agent-Based Technology*. IntechOpen, 2012.
- Niko Speybroeck, Carine Van Malderen, Sam Harper, Birgit Müller, and Brecht Devleesschauwer. Simulation models for socioeconomic inequalities in health: a systematic review. *International journal of environmental research and public health*, 10(11):5750–5780, 2013.
- Lei Wang, Jingsen Zhang, Hao Yang, Zhiyuan Chen, Jiakai Tang, Zeyu Zhang, Xu Chen, Yankai Lin, Ruihua Song, Wayne Xin Zhao, et al. User behavior simulation with large language model based agents. *arXiv preprint arXiv:2306.02552*, 2023.
- Zengqing Wu, Run Peng, Xu Han, Shuyuan Zheng, Yixin Zhang, and Chuan Xiao. Smart agent-based modeling: On the use of large language models in computer simulations. *arXiv preprint arXiv:2311.06330*, 2023.
- Ziyi Yang, Zaibin Zhang, Zirui Zheng, Yuxian Jiang, Ziyue Gan, Zhiyu Wang, Zijian Ling, Jinsong Chen, Martz Ma, Bowen Dong, et al. Oasis: Open agents social interaction simulations on one million agents. *arXiv preprint arXiv:2411.11581*, 2024.
- Wei Zhang, Andrea Valencia, and Ni-Bin Chang. Synergistic integration between machine learning and agent-based modeling: A multidisciplinary review. *IEEE Transactions on Neural Networks and Learning Systems*, 34(5):2170–2190, 2021.
- Weitong Zhang, Chengqi Zang, and Bernhard Kainz. Truth or deceit? a bayesian decoding game enhances consistency and reliability. *arXiv preprint arXiv:2410.01064*, 2024.
- Xuhui Zhou, Hao Zhu, Leena Mathur, Ruohong Zhang, Haofei Yu, Zhengyang Qi, Louis-Philippe Morency, Yonatan Bisk, Daniel Fried, Graham Neubig, et al. Sotopia: Interactive evaluation for social intelligence in language agents. *arXiv preprint arXiv:2310.11667*, 2023.

A APPENDIX

A.1 REPRESENTATIVE ABM: PROOF FOR PROPOSITION 2.1.

In a two-bidder Dutch auction where valuations are drawn independently from $U[0,1]$ and bidders have exponential utility functions $u(x) = e^{rx}$ with heterogeneous risk parameters, the equilibrium bidding strategy is $\beta(v, r) = v - \frac{1}{r}$.

Consider a Dutch auction where the price p starts high and decreases continuously. At each price p , a bidder must decide whether to accept the current price. Let $\beta(v, r)$ denote the price at which a bidder with value v and risk parameter r plans to accept. At the optimal acceptance price b , the expected utility of accepting must equal the marginal expected utility of waiting:

$$\left. \frac{d}{dp} \right|_{p=b} E[u(v-p)|p \leq b] = 0$$

Let $F(\beta^{-1}(p))$ denote the probability that the other bidder has not yet accepted at price p . Then the expected utility is:

$$E[u(v-p)|p \leq b] = F(\beta^{-1}(b))e^{r(v-b)}$$

Taking the derivative with respect to b :

$$\frac{d}{db} [F(\beta^{-1}(b))e^{r(v-b)}] = (\beta^{-1})'(b)e^{r(v-b)} - F(\beta^{-1}(b))re^{r(v-b)} = 0$$

Conjecture a linear bidding strategy of the form:

$$\beta(v, r) = v - c(r)$$

Therefore:

$$\beta^{-1}(b) = b + c(r)$$

$$(\beta^{-1})'(b) = 1$$

Substituting into the FOC:

$$e^{r(v-b)} - (b + c(r))re^{r(v-b)} = 0$$

$$1 - (b + c(r))r = 0$$

$$b + c(r) = \frac{1}{r}$$

At equilibrium, $b = \beta(v, r) = v - c(r)$, so:

$$v - c(r) + c(r) = \frac{1}{r}$$

$$v = \frac{1}{r}$$

Therefore:

$$c(r) = \frac{1}{r}$$

And the equilibrium bidding strategy is:

$$\beta(v, r) = v - \frac{1}{r}$$

To verify this is indeed an equilibrium, since the strategy is strictly increasing in v , ensuring higher-value bidders accept earlier. No bidder has incentive to deviate as accepting earlier reduces payoff if won, while accepting later reduces probability of winning.

A.2 BiD: SIMULTANEOUS COMMUNICATION

(Existence of Perfect Bayesian Equilibrium) A mixed strategy PBE equilibrium exists for the last round of MRMP-DA with Simultaneous Reporting.

Let Σ denote the space of probability measures over $[0, 1] \times [0, 1]$, representing mixed strategies over prices and reports, endowed with the weak* topology. The expected payoff for an agent using mixed strategy y against opponent's strategy σ is:

$$E[\pi(y, \sigma)] = \int \int \int [c \cdot \mathbf{1}\{b_i > b_j\} \cdot e^{f(v)(v-b_i)} + (1-c) \cdot \int_{\hat{v}_j}^1 \mathbf{1}\{b_i > b_j\} \cdot e^{f(v)(v-b_i)} dw] d\sigma(b_j, \hat{v}_j) dy(b_i, \hat{v}_i)$$

We establish the following properties:

Lemma A.1 (Strategy Space) Σ is nonempty, convex, and compact in the weak* topology.³

Nonemptiness follows directly from the existence of pure strategies. Convexity follows directly from the definition of probability measures. First, we show that Σ is a subset of the dual space of continuous functions on $[0, 1] \times [0, 1]$.

As for the compactness, let $C([0, 1] \times [0, 1])$ be the space of continuous functions on $[0, 1] \times [0, 1]$ with supremum norm.

By definition, any $\sigma \in \Sigma$ is a probability measure, so for any $f \in C([0, 1] \times [0, 1])$:

$$\left| \int f d\sigma \right| \leq \|f\|_\infty$$

Therefore, Σ is contained in the unit ball of the dual space $C([0, 1] \times [0, 1])^*$.

Then we need to show that Σ is closed in the weak* topology. To see this, let $\{\sigma_n\}$ be a sequence in Σ converging to σ in the weak* topology. Then first, for any non-negative $f \in C([0, 1] \times [0, 1])$: $\int f d\sigma \geq 0$ by weak* convergence, and second, $\int 1 d\sigma = \lim_{n \rightarrow \infty} \int 1 d\sigma_n = 1$. Therefore we get that σ is a probability measure, so $\sigma \in \Sigma$.

By the Banach-Alaoglu theorem, the unit ball in $C([0, 1] \times [0, 1])^*$ is compact in the weak* topology. Since Σ is a closed subset of a compact set in the weak* topology, it is compact in the weak* topology.

To establish sequential compactness, let $\{\sigma_n\}$ be a sequence in Σ . By weak* compactness, there exists a subsequence $\{\sigma_{n_t}\}$ and a measure $\sigma \in C([0, 1] \times [0, 1])^*$ such that for all $f \in C([0, 1] \times [0, 1])$:

$$\lim_{t \rightarrow \infty} \int f d\sigma_{n_t} = \int f d\sigma$$

By the closedness shown above, $\sigma \in \Sigma$.

Lemma A.2 (Payoff Continuity) $E[\pi(y, \sigma)]$ is continuous in the weak* topology.

Consider sequences $y_n \rightarrow y$ and $\sigma_n \rightarrow \sigma$ in the weak* topology. The integrand

$$h(b_i, \hat{v}_i, b_j, \hat{v}_j) = c \cdot \mathbf{1}\{b_i > b_j\} \cdot e^{f(v)(v-b_i)} + (1-c) \cdot \int_{\hat{v}_j}^1 \mathbf{1}\{b_i > b_j\} \cdot e^{f(v)(v-b_i)} dw$$

is bounded since $e^{f(v)(v-b_i)}$ is bounded on $[0, 1] \times [0, 1]$. The discontinuities in the indicator function occur only on the measure-zero set $\{b_i = b_j\}$. By the bounded convergence theorem $|E[\pi(y_n, \sigma_n)] - E[\pi(y, \sigma)]| \rightarrow 0$ as $n \rightarrow \infty$

³the weak* topology on the space of probability measures Σ over $[0, 1] \times [0, 1]$ is defined by convergence against continuous bounded functions. Specifically, a sequence of measures $\{\sigma_n\}$ converges to σ if: $\int f d\sigma_n \rightarrow \int f d\sigma$ for all continuous bounded functions f

Lemma A.3 (Upper Hemicontinuity) *The best response correspondence $BR : \Sigma \rightarrow 2^\Sigma$ is upper hemicontinuous.*

Take any alternative strategy $z \in \Sigma$ (a probability measure over $[0, 1] \times [0, 1]$). Since $y_n \in BR(\sigma_n)$, for each n :

$$\begin{aligned} E[\pi(y_n, \sigma_n)] &= \int \int \int [c \cdot \mathbf{1}\{p_i > p_j\} \cdot e^{f(v)(v-p_i)} + \\ &\quad (1-c) \cdot \int_{R_j}^1 \mathbf{1}\{p_i > p_j\} \cdot e^{f(v)(v-p_i)} dw] d\sigma_n(p_j, R_j) dy_n(p_i, R_i) \\ &\geq E[\pi(z, \sigma_n)] \end{aligned}$$

Consider $|E[\pi(y_n, \sigma_n)] - E[\pi(y, \sigma)]|$. This difference converges to zero because:

- The integrand is bounded by some M (due to the boundedness of the exponential term on $[0, 1]$)
- The indicator function's discontinuities occur on measure-zero set $\{p_i = p_j\}$
- $\sigma_n \rightarrow \sigma$ and $y_n \rightarrow y$ in weak* topology

Therefore, $\exists N_1$ such that for $n > N_1$: $|E[\pi(y_n, \sigma_n)] - E[\pi(y, \sigma)]| < \varepsilon/2$

Similarly, $|E[\pi(z, \sigma_n)] - E[\pi(z, \sigma)]| \rightarrow 0$. Therefore $\exists N_2$ such that for $n > N_2$: $|E[\pi(z, \sigma_n)] - E[\pi(z, \sigma)]| < \varepsilon/2$

Let $N = \max\{N_1, N_2\}$. For $n > N$:

$$E[\pi(y, \sigma)] > E[\pi(y_n, \sigma_n)] - \varepsilon/2 \geq E[\pi(z, \sigma_n)] - \varepsilon/2 > E[\pi(z, \sigma)] - \varepsilon$$

Since ε, z was arbitrary, $E[\pi(y, \sigma)] \geq E[\pi(z, \sigma)]$ and $y \in BR(\sigma)$

The strategy space Σ is nonempty, convex, and compact by Lemma A.1 The payoff function is continuous by Lemma A.2, and the best response correspondence is upper hemicontinuous with nonempty, convex values by Lemma A.3. Therefore, by Glicksberg's fixed point theorem Glicksberg (1952), there exists a mixed strategy σ^* such that $\sigma^* \in BR(\sigma^*)$, establishing the existence of a mixed strategy equilibrium.

A.3 BID: SEQUENTIAL COMMUNICATION

Proof for Theorem 2.3

Theorem A.4 (Existence of Perfect Bayesian Equilibrium) *A Perfect Bayesian Equilibrium exist for the last round of MRMP-DA with Sequential Reporting.*

The proof is rather similar to the one with the simultaneous reporting scheme; the only difference comes from low valuation's reporting strategy such that it now maps from her own valuation and opponent's report to an optimal report.

1) First, we define the strategy spaces. Let Σ_i be the space of probability measures over $[0, 1] \times [0, 1]$ representing mixed strategies over reports and bids, endowed with the weak* topology. For $\sigma_i \in \Sigma_i$:

- $\sigma_i^R : [0, 1] \rightarrow \Delta([0, 1])$ maps types to distributions over reports
- $\sigma_i^B : [0, 1] \times [0, 1] \times [0, 1] \rightarrow \Delta([0, 1])$ maps (type, own report, opponent's report) to distributions over bids

2) For any observed report \hat{v}_H , the posterior belief $\mu(\cdot | \hat{v}_H)$ is:

$$\mu(v_H | \hat{v}_H) = c \cdot \frac{\sigma_H^R(\hat{v}_H | v_H) f(v_H)}{\int \sigma_H^R(\hat{v}_H | v) f(v) dv} + (1-c) \cdot U[\hat{v}_H, 1]$$

where f is the uniform density on $[0, 1]$.

3) The low valuation bidder's best response problem is:

$$BR_L(v_L, \hat{v}_H) = \arg \max_{\sigma_L \in \Sigma_L} \int_{\hat{v}_H}^1 \int_0^{v_L} \int_0^1 \pi_L(\hat{v}_L, b_L, \hat{v}_H, b_H, v_L, v_H) d\sigma_L^B(b_L|v_L, \hat{v}_L, \hat{v}_H) d\sigma_L^R(\hat{v}_L|v_L) d\mu(v_H|\hat{v}_H)$$

subject to $\hat{v}_L \leq v_L$ (no over-reporting constraint).

4) The high valuation bidder's best response problem is:

$$BR_H(v_H) = \arg \max_{\sigma_H \in \Sigma_H} \int_0^{v_H} \int_0^{v_H} \int_0^1 \pi_H(\hat{v}_H, b_H, \hat{v}_L, b_L, v_H, v_L) d\sigma_H^B(b_H|v_H, \hat{v}_H, \hat{v}_L) d\sigma_H^R(\hat{v}_H|v_H) dF(v_L)$$

subject to $\hat{v}_H \leq v_H$.

5) Define correspondence $\Phi : \Sigma_H \times \Sigma_L \times \mathcal{P}(V) \rightarrow 2^{\Sigma_H \times \Sigma_L \times \mathcal{P}(V)}$ by:

$$\Phi(\sigma_H, \sigma_L, \mu) = (BR_H(\sigma_L, \mu), BR_L(\sigma_H, \mu), \text{Bayes}(\sigma_H, \sigma_L))$$

6) Φ is nonempty, convex and compact with the same argument from theorem ??.

Lemma A.5 (Closed Graph Property) *The correspondence $\Phi : \Sigma_H \times \Sigma_L \times \mathcal{P}(V) \rightarrow 2^{\Sigma_H \times \Sigma_L \times \mathcal{P}(V)}$ has a closed graph in the product topology.*

Take sequences:

$$\begin{aligned} (\sigma_H^n, \sigma_L^n, \mu^n) &\rightarrow (\sigma_H, \sigma_L, \mu) \text{ in the product topology} \\ (\sigma_H'^n, \sigma_L'^n, \mu'^n) &\in \Phi(\sigma_H^n, \sigma_L^n, \mu^n) \end{aligned}$$

with

$$(\sigma_H'^n, \sigma_L'^n, \mu'^n) \rightarrow (\sigma_H', \sigma_L', \mu') \text{ in the product topology}$$

We need to show that $(\sigma_H', \sigma_L', \mu') \in \Phi(\sigma_H, \sigma_L, \mu)$.

First, consider the best response correspondence for the high valuation bidder. Since $\sigma_H'^n \in BR_H(\sigma_L'^n, \mu'^n)$, we have for all alternative strategies $\tilde{\sigma}_H \in \Sigma_H$:

$$\begin{aligned} &\int_0^{v_H} \int_0^{v_H} \int_0^1 \pi_H(\hat{v}_H, b_H, \hat{v}_L, b_L, v_H, v_L) d\sigma_H'^n(b_H|v_H, \hat{v}_H, \hat{v}_L) d\sigma_H'^n(\hat{v}_H|v_H) dF(v_L) \\ &\geq \int_0^{v_H} \int_0^{v_H} \int_0^1 \pi_H(\hat{v}_H, b_H, \hat{v}_L, b_L, v_H, v_L) d\tilde{\sigma}_H(b_H|v_H, \hat{v}_H, \hat{v}_L) d\tilde{\sigma}_H(\hat{v}_H|v_H) dF(v_L) \end{aligned}$$

This inequality converges because:

1) The payoff function π_H can be decomposed as:

$$\pi_H = e^{f(v_H)(v_H - b_H)} \cdot \mathbf{1}\{b_H > b_L\}$$

2) The exponential term $e^{f(v_H)(v_H - b_H)}$ is bounded on $[0, 1] \times [0, 1]$ by some constant M since $f(v_H)$ is continuous.

3) The indicator function $\mathbf{1}\{b_H > b_L\}$ creates discontinuities only on the measure-zero set $\{b_H = b_L\}$.

4) By the bounded convergence theorem Royden & Fitzpatrick (2010):

$$\begin{aligned} &\lim_{n \rightarrow \infty} \int_0^{v_H} \int_0^{v_H} \int_0^1 \pi_H d\sigma_H'^n d\sigma_H'^n dF \\ &= \int_0^{v_H} \int_0^{v_H} \int_0^1 \pi_H d\sigma_H' d\sigma_H' dF \end{aligned}$$

Therefore, in the limit:

$$\sigma'_H \in BR_H(\sigma_L, \mu)$$

Similar argument holds for the low valuation's best response.

For the belief component, note that with Bayes updating, we have

$$\mu^n(v_H|\hat{v}_H) = c \cdot \frac{\sigma_H^n(\hat{v}_H|v_H)f(v_H)}{\int \sigma_H^n(\hat{v}_H|v)f(v)dv} + (1-c) \cdot U[\hat{v}_H, 1]$$

As $n \rightarrow \infty$, the first term converges by weak* convergence of σ_H^n and continuity of the ratio and the second term is constant in n thus automatically converges.

Thus, $(\sigma'_H, \sigma'_L, \mu') \in \Phi(\sigma_H, \sigma_L, \mu)$, establishing that Φ has a closed graph.

Best responses exist by compactness of strategy spaces, convexity follows from linearity of expected utility, and Bayes update is single-valued, thus Values of Φ are nonempty and convex.

7) By Kakutani's fixed point theorem:

$$\exists(\sigma_H^*, \sigma_L^*, \mu^*) \in \Phi(\sigma_H^*, \sigma_L^*, \mu^*)$$

This fixed point constitutes a PBE because:

- Strategies are mutual best responses given beliefs
- Beliefs are consistent with Bayes' rule where possible
- Sequential rationality holds at each information set

Then, we show that the equilibrium utility is weakly higher than the Single-Round One-Phase DA case by showing that our Single-Round Two-Phase DA has a larger strategy space.

To simulate the Single-Round One-Phase case, we can let bidder 1 **randomize reporting uniformly over** $[0, 1]$ and bidder 2 **ignore the reporting and play the Single-Round One-Phase equilibrium strategy**, therefore, we know that both bidders can be weakly better off because the strategy space is strictly larger.