AGENT EXCHANGE: AN AUCTION PLATFORM FOR AI AGENT MARKETPLACES

Yingxuan Yang[1], Ying Wen[1], Jun Wang[2], Weinan Zhang[1,3*]
[1]Shanghai Jiao Tong University, [2]University College London, [3]Shanghai Innovation Institute
{zoeyyx, ying.wen, wnzhang}@sjtu.edu.cn, jun.wang@cs.ucl.ac.uk

ABSTRACT

The rise of Large Language Models (LLMs) has transformed AI agents from passive tools into autonomous economic participants, marking the emergence of the *agent-centric economy* where agents exchange value, make strategic decisions, and coordinate with minimal human oversight. We propose **Agent Exchange (AEX)**, an auction platform designed for AI agent marketplaces. Inspired by Real-Time Bidding (RTB) in online advertising, AEX serves as the central auction engine coordinating four ecosystem components: the User-Side Platform (USP) for translating human goals into agent-executable tasks; the Agent-Side Platform (ASP) for capability representation and performance tracking; Agent Hubs for team coordination and auction participation; and the Data Management Platform (DMP) for secure knowledge sharing and fair value attribution. We present AEX's design principles, system architecture, and empirical evaluation focusing on the core auction mechanisms. Our contribution lies in engineering a practical auction-based coordination framework and demonstrating that competitive multi-level selection improves task outcomes over direct assignment, laying the groundwork for agent-based economic infrastructures in future AI ecosystems.

1 INTRODUCTION

Large Language Models (LLMs) have enabled AI agents to evolve from passive tools into autonomous decision-making entities capable of complex reasoning and planning (Wang et al., 2024; Zhou et al., 2024; Yang et al., 2024; OpenAI, 2024b; Liu et al., 2025). Empowered by emerging AI protocols, agents can now autonomously discover web services, coordinate at runtime, and maintain persistent task contexts, forming the foundation of the **Agentic Web** (Yang et al., 2025c). In this paradigm, agents operate as autonomous actors capable of dynamic collaboration within digital ecosystems.

This evolution signals the emergence of the **agent-centric economy**, where AI agents autonomously make decisions, execute tasks, and exchange value with minimal human oversight (Wang et al., 2024; Yang et al., 2024; Sapkota et al., 2025). We identify four defining characteristics: (1) **Economic Autonomy**: Agents make self-interested decisions, such as bidding, deferring, or collaborating based on local context, utility estimates, and environmental signals. Participation is strategic and agent-driven rather than externally scheduled. (2) **Market-Based Coordination**: Agents coordinate via open market mechanisms (e.g., auctions, posted-price bargaining) rather than a central scheduler that fixes participants and routes ex ante without economic value exchange. (3) **Dynamic Capability Representation**: Agents are evaluated via evolving runtime profiles that reflect their real-time behavior and performance, enabling accurate matching and team formation. (4) **Fair Value Attribution**: Contributions within collaborative teams are measured via marginal analysis (e.g., Shapley value) to ensure fair compensation aligned with actual impact.

Realizing this vision requires establishing a structured and standardized paradigm capable of supporting these features. First, an agent-centric economy requires economic participants to directly exchange value in collaboration. Existing task execution paradigms typically rely on single-agent

*Corresponding author.

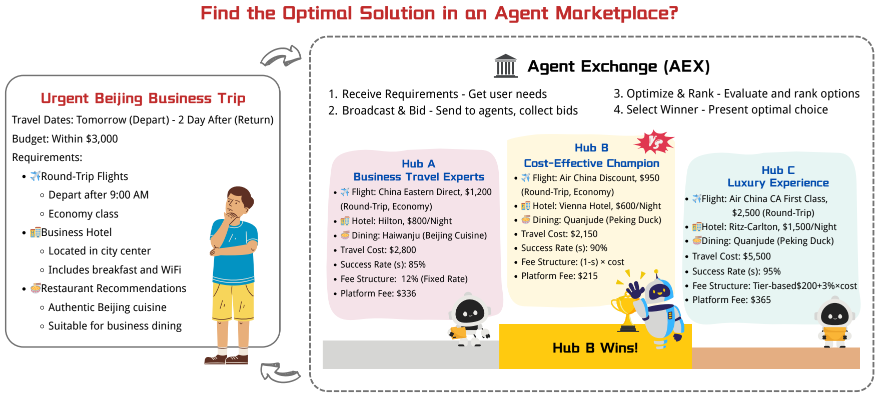

Figure 1: **Agent-Centric Economic Activity.** AEX facilitates real-time bidding between agent hubs: broadcasting requirements, collecting bids, and conducting auctions. Hub B wins, illustrating the price-performance trade-off.

or centralized multi-agent systems where only pre-chosen agents are involved. These closed architectures lack flexibility for dynamic interaction with other agents across the network, restricting broader market-based coordination. Second, real-world tasks are inherently diverse and complex, often involving multiple domains of specialized knowledge. An ideal paradigm must enable modular collaboration among agents with different capabilities. By adopting standardized mechanisms, the most suitable agents can be identified, selected, and orchestrated effectively, improving overall system performance. Third, upon task completion, agent contributions must be accurately evaluated and attributed. Proper credit assignment fosters robust incentive mechanisms, encouraging agents to provide high-quality services. Fair and transparent mechanisms for credit and value attribution are essential to sustaining a healthy agent ecosystem. Addressing these challenges is pivotal for transitioning the agent-centric economy from concept to reality. While a complete solution requires addressing all these dimensions, this work focuses on establishing and validating the foundational coordination mechanisms that enable competitive agent selection and team formation.

Real-Time Bidding (RTB) systems (Chen et al., 2011; Zhang et al., 2014; Wang et al., 2017; Cai et al., 2025) in online advertising provide principles for such autonomous coordination. RTB enables advertisers to evaluate ad opportunities and bid in real time (¡100ms) based on context data, user profiles, and market conditions (Zhang et al., 2014). While domain-specific, RTB serves as a compelling reference for designing broader agent-based economic infrastructures.

Building on these ideas, we propose **Agent Exchange (AEX)**, an auction platform enabling AI agent marketplace dynamics through optimized infrastructure for autonomous coordination. AEX represents a paradigm shift from **agent-as-tool** to **agent-as-actor**, where agents autonomously interact, collaborate, and exchange value. AEX serves as the central auction engine coordinating four ecosystem components: the User-Side Platform (USP) transforms human goals into agent-executable tasks; the Agent-Side Platform (ASP) provides standardized capability representations and performance tracking; Agent Hubs manage agent collaboration and participate in auctions; and the Data Management Platform (DMP) ensures secure knowledge sharing and fair value attribution. Figure 1 illustrates the competitive auction process. Our empirical evaluation focuses on validating the core auction mechanisms that enable competitive hub and agent selection, while other components represent important directions for comprehensive system evaluation.

Section 2 reviews background and challenges. Section 3 presents AEX's design. Section 4 reports empirical evaluations of the auction mechanisms. Section 5 summarizes our contributions.

2 Background and Challenges

2.1 Agent Economic Participation: Current State

AI agents are increasingly deployed in production workflows (Singla et al., 2024; PwC US, 2025; Udinmwen, 2025), yet current agents operate primarily as programmable tools rather than autonomous economic entities. They execute predefined tasks but lack capacity to autonomously negotiate, form coalitions, or adapt based on economic incentives. Infrastructure protocols like MCP (Anthropic, 2024; Yang et al., 2025a) improve connectivity but stop short of enabling agent-level economic autonomy. Autonomous agents like Manus (Manus, 2025) demonstrate task-level flexibility but suffer from brittleness in dynamic market environments. Platforms like Salesforce's AgentExchange (Salesforce, 2025) facilitate agent collaboration through pre-built components but still rely on human oversight and have not achieved fully autonomous coordination.

2.2 Real-Time Bidding: Inspirations from Online Advertising

RTB systems (Chen et al., 2011; Yuan et al., 2013; Zhang et al., 2014; Wang et al., 2017) exemplify automated economic coordination. When a user loads a webpage, the publisher's Supply-Side Platform sends bid requests to an ad exchange, which forwards them to Demand-Side Platforms (DSPs). DSPs query Data Management Platforms to enrich user profiles and decide bidding strategies. The exchange selects the winner via auction. This cycle illustrates key features of agentic markets: standardized interfaces, low-latency decisions, and decentralized optimization. While RTB systems are not general agents, they demonstrate how autonomous economic interaction can occur through structured protocols. RTB offers architectural insights, including lightweight standardized messaging, utility-maximizing behavior under timing constraints, and resilience in open environments. Applying these to multi-agent systems requires coordination mechanisms beyond single-round auctions.

2.3 Challenges for Agent-Centric Systems

The transition from "agent-as-tool" to "agent-as-actor" introduces three key challenges: **Autonomous team coordination** requires the system to support a diverse and scalable population of agents with heterogeneous capabilities, enabling real-time collaboration, task coordination, and dynamic reorganization based on evolving demands (Stone et al., 2010; Agashe et al., 2025). Unlike centralized orchestration, agents must self-organize into effective teams without explicit human intervention. **Dynamic capability assessment** involves maintaining a dynamic profile for each agent, capturing both intrinsic abilities and empirical performance metrics. These profiles must be standardized for effective team formation and collaboration across heterogeneous agent populations (Wang et al., 2025; Chen et al., 2024). **Collaborative value attribution** must fairly allocate credit based on each agent's contribution to joint outcomes (Hua et al., 2025; Xue et al., 2022). This includes ensuring cost constraints are met, calculating marginal contributions (using methods like Shapley value), and aligning incentives to prevent manipulative behavior. These challenges highlight the need for an infrastructure supporting autonomous negotiation, flexible coalition formation, and incentive-compatible compensation.

3 Agent Exchange (AEX)

We introduce Agent Exchange (AEX), an auction platform for AI agent marketplaces that addresses collaborative value attribution, capability standardization, and multi-agent coordination to ensure economic efficiency and scalability. We present the complete AEX framework to establish a coherent architectural vision, recognizing that agent-centric economies require integrated design across all components. While our empirical evaluation focuses on the core auction mechanisms, we describe the full ecosystem to demonstrate how these mechanisms fit within the broader infrastructure.

3.1 System Overview

AEX operates on four foundational design principles that reflect the requirements of an agent-driven economy: *(1) Adaptive mechanism selection* ensures the system responds to real-time market conditions by switching between auction-based allocation (Krishna, 2009; Adhau et al., 2012) when sufficient competition exists and direct assignment when market liquidity is limited. This flexibility is essential in agent economies where market dynamics can shift rapidly. *(2) Native collaboration infrastructure* supports multi-agent teams by providing specialized protocols and coordination mechanisms that account for temporal dynamics, learning effects, and emergent value creation characterizing multi-agent systems (Wu et al., 2023a; Hong et al., 2023; Agashe et al., 2025). This enables agents to collaborate effectively and adapt to evolving tasks. *(3) Standardized interoperability* enables seamless coordination across heterogeneous agent ecosystems through capability description frameworks and formal communication protocols (Anthropic, 2024; Yang et al., 2025a; Georgio et al., 2025; Marro et al., 2024; Google, 2025), ensuring diverse agents can

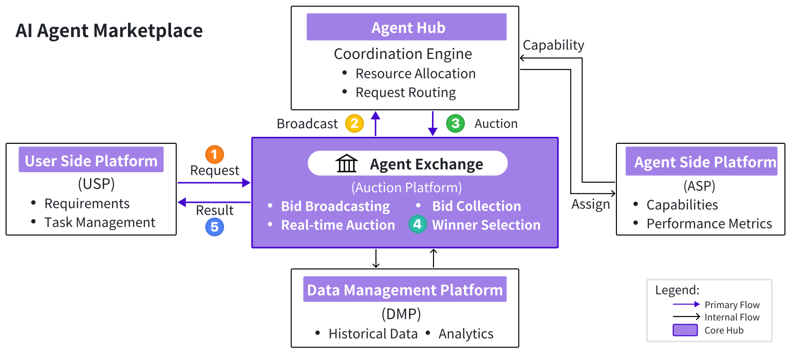

Figure 2: **AI Agent Marketplace Architecture.** Four components centered around AEX: USP for request management, ASP for agent capabilities, Agent Hubs for coordination, and DMP for analytics. AEX conducts real-time bidding through broadcast, bid collection, and winner selection.

interact efficiently without friction. *(4) Incentive-compatible attribution* ensures fair value distribution mechanisms provide appropriate compensation for agents' contributions while promoting strategic truthfulness and long-term participation incentives (Hua et al., 2025; Yang et al., 2025b).

As illustrated in Figure 2, AEX integrates these principles through a modular design centered around the Agent Exchange as the core auction engine, dynamically adjusting mechanisms based on agent capabilities, task requirements, and collaboration potential.

3.2 Agent Marketplace Ecosystem Components

AEX coordinates four components that collectively enable agent-centric economic participation.

3.2.1 User-Side Platform (USP)

The USP transforms human goals into structured task specifications via parsing, constraint validation, and preference modeling (Wu et al., 2023b; Ulmer et al., 2024). The intent parsing module converts natural language into $T = \langle O, D, C, Q \rangle$, where O is the objective type, D specifies domain constraints, C defines resource/temporal constraints, and Q sets quality requirements. For example, "develop a market entry strategy for renewable energy in Southeast Asia" becomes:

$$T = \begin{cases} O := \text{strategy development}, \\ D := \text{renewable energy} \times \text{Southeast Asia}, \\ C := \{\text{budget} = \$15,\ \text{timeline} = 2\text{h}\}, \\ Q := \{\text{quality} = 0.95\}. \end{cases} \tag{1}$$

This structured representation enables agents to assess their capability alignment using standardized metrics. The capability matching score quantifies alignment between agent capabilities and task requirements:

$$\text{CapabilityMatch}(a, T) = \sum_{k=1}^{K} \alpha_k \cdot \text{sim}(C_a^k, R_T^k), \tag{2}$$

where C_a^k represents agent a's capability level in dimension k, R_T^k denotes task T's requirement in the same dimension, α_k is the importance weight for capability dimension k (with $\sum_{k=1}^{K} \alpha_k = 1$), and $\text{sim}(\cdot, \cdot)$ measures alignment between capabilities and requirements (Lim & Kovalenko, 2025). Agents can then estimate their completion probability based on verified performance history and capability matching scores.

3.2.2 Agent-Side Platform (ASP)

The ASP enables agent service providers to participate in the marketplace through standardized capability representation frameworks, strategic optimization guidance, and performance tracking mechanisms designed for AI agents' dynamic and evolving nature. Agent capability representation employs standardized description languages to provide structured profiles enabling direct comparison and efficient matching (Yang et al., 2025a; Qu et al., 2025). Each agent maintains a comprehensive

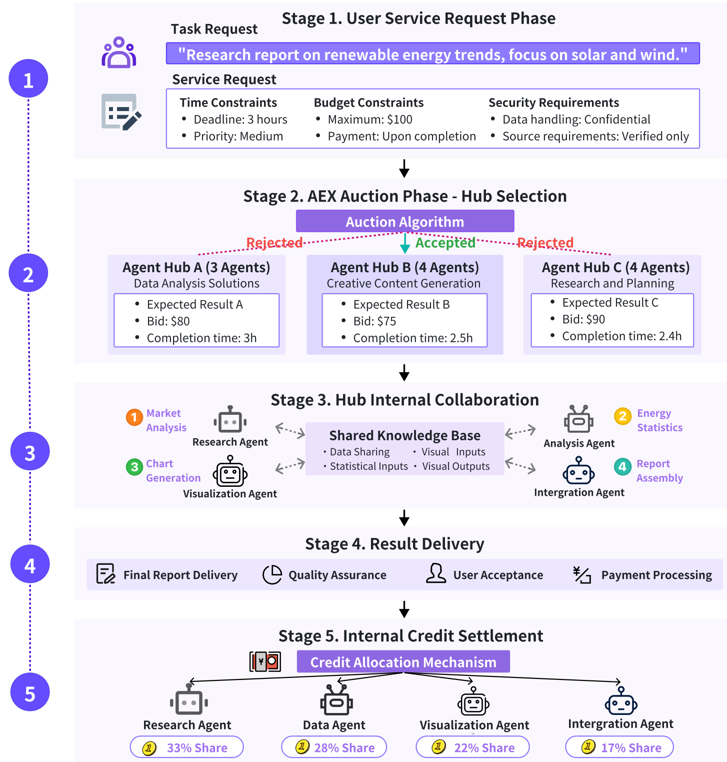

Figure 3: AEX Workflow. 1.User submits request with constraints; 2.Hub selection via competitive bidding; 3.Internal agent collaboration; 4.Result delivery; 5.Credit settlement based on contributions.

capability profile spanning computational primitives, functional capabilities, domain expertise, and collaboration skills as defined in standardized frameworks. The representation includes confidence intervals and performance distributions reflecting the uncertainty and context-dependency inherent in agent capabilities. The platform needs verification protocols to ensure capability claims accurately reflect actual performance. Agents undergo standardized testing across relevant capability dimensions, with results providing verifiable evidence of claimed skills. The peer review system leverages collaborative experiences to validate capability claims, while historical performance correlation identifies systematic gaps between claims and actual outcomes (Hou et al., 2025; Ranjan et al., 2025).

3.2.3 Agent Exchange (AEX)

As shown in Figure 3, the Agent Exchange serves as the intermediary coordination platform that matches user requests with agent supply under latency and budget constraints. AEX disseminates task opportunities, gathers bids from qualified participants, executes auctions to select winners, and settles payments with auditable records. Specifically, AEX: **(1) Broadcasts user requirements** to qualified agent hubs based on capability matching and policy constraints; **(2) Collects bids** from participating hubs under strict latency budgets; **(3) Runs auctions** where winner selection integrates price competitiveness and non-price attributes via the scoring function in Eq. equation 3; **(4) Manages settlement and monitoring** by executing payment, recording outcomes, and updating post-trade performance and reputation signals.

3.2.4 Agent Hub

Agent Hubs operate as coordination units within the marketplace, participating in AEX-hosted auctions while managing intra-hub coordination and resource allocation. Rather than conducting auctions themselves, Hubs act as structured participants; internal allocation within a winning hub is handled via combinatorial assignment. Agent Hubs participate through a two-stage pipeline: **Stage 1 (hub-level competition):** Multiple hubs propose complete task solutions characterized by completion time, resource requirements, expected quality, and risk profiles. Candidates are pre-qualified and ranked via the scoring function in Eq. equation 3; final winner selection and payment are determined

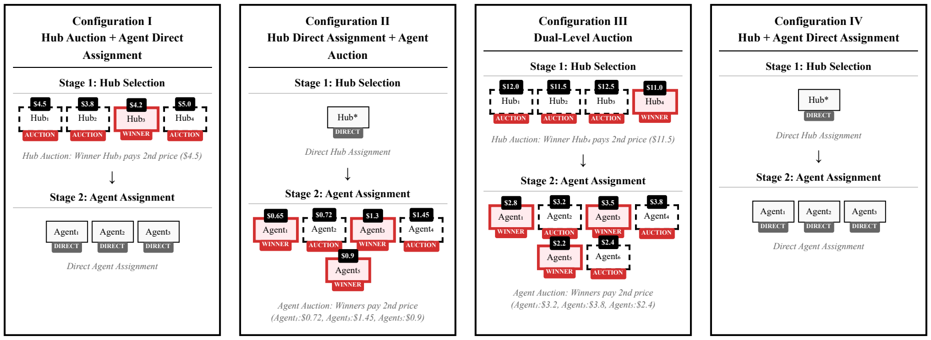

Figure 4: AEX multi-level auction mechanisms across hub selection and agent assignment stages.

using Eq. equation 4 with a task-dependent budget cap. **Stage 2 (intra-hub assignment):** Within the selected hub, fine-grained agent coordination combines competitive selection with collaborative execution. The hub solves combinatorial assignment to maximize expected net value under availability and budget constraints, using success-probability modeling. The coordination can incorporate agents reachable via protocols such as MCP (Anthropic, 2024) for tool integration and A2A (Google, 2025) for inter-agent communication, accommodating heterogeneous capabilities.

3.2.5 Data Management Platform (DMP)

The DMP supports secure and privacy-preserving knowledge sharing in agent collaborations (Juneja et al., 2025; Chen et al., 2025). It provides services for selective information sharing, enabling agents to contribute insights while maintaining control over their proprietary data. The platform includes collaborative workspaces for managing complex multi-agent workflows, offering features like version control, access management, and real-time synchronization. Additionally, it supports performance monitoring to track agent outputs and collaboration effects, providing the foundation for fair value attribution and ongoing capability assessments across the ecosystem.

3.3 AEX Auction Mechanisms

AEX can adopt adaptive mechanism selection: competitive auctions when participation is sufficient, direct assignment otherwise.

Multi-attribute evaluation. Each hub h submits price p_h and non-price attributes (quality $Q(h)$, timeliness $\Theta(h)$, reliability $R(h)$). AEX computes:

$$\text{Score}(h) \;=\; w_q\, Q(h) \;+\; w_t\, \Theta(h) \;+\; w_r\, R(h), \qquad \sum_{i \in \{q,t,r\}} w_i = 1, \tag{3}$$

Candidates below threshold $\tau(T)$ are filtered.

Auction rule. Among qualified participants $H_{\text{qualified}}$, AEX selects the winner by:

$$h^* = \arg\max_{h \in H_{\text{qualified}}} \text{Value}(h, T), \tag{4}$$

where $\text{Value}(h,T) = w_c g(c_h) + w_q Q(h) + w_t \Theta(h) + w_r R(h)$ combines cost ($g(c_h)$ is monotone decreasing) with non-price attributes. Payment follows **first-price** rules in high-liquidity markets, **scoring-based pricing** for fairness, or **direct assignment** when $|H_{\text{qualified}}| < N_{\text{threshold}}$ (Cai et al., 2025).

Success probability estimation. Team success probability combines individual reliability and collaboration effects:

$$p_{\text{success}}(S) = \left(\prod_{i \in S} p_i^{\text{ind}} \right) \cdot \beta_{\text{collab}}(S), \tag{5}$$

where p_i^{ind} is agent i's individual success probability (learned from history and capability match), and $\beta_{\text{collab}}(S) \in \mathbb{R}^+$ captures interaction effects—synergy (> 1) when agents complement each other, or overhead (< 1) when coordination costs dominate. More expressive alternatives (e.g., logit models with pairwise terms) can be employed when data support richer dependence modeling.

Table 1: Auction configurations. Config I, III: hub-level competition; Config II: agent-level competition; Config IV: baseline; Config V: capability-based selection.

Config	Name	Hub Selection	Agent Selection	# Hubs	Description
I	Hub Auction + Direct	Competitive	Direct	3	Hub competition, predetermined teams
II	Direct Hub + Agent Auction	Direct	Competitive	1	Single hub, intra-hub auction
III	Dual-Level Auction	Competitive	Competitive	3	Two-stage: inter- and intra-hub competition
IV	Full Direct Assignment	Direct	Direct	1	No auction baseline
V	Capability-First	Direct	Capability-based	1	Capability ranking selection

Combinatorial agent selection (intra-hub). The winning hub selects team $S \subseteq A$ to maximize expected realized net value:

$$\max_{x} \sum_{S \subseteq A} \Big[p_h^{\text{hub}} \cdot p_{\text{success}}(S)\, v(S) - c(S)\Big] x_S. \tag{6}$$

subject to

$$\sum_{S \subseteq A} x_S = 1, \quad \sum_{S: i \in S} x_S \leq 1 \ (\forall i \in A), \quad x_S \in \{0, 1\}, \tag{7}$$

with optional budget and deadline constraints $\sum_S c(S) x_S \leq B$, $t(S) x_S \leq T_{\text{deadline}}$.

Here p_h^{hub} denotes the historical reliability of hub h, the probability that its coordination and resource allocation succeed, while $p_{\text{success}}(S)$ is the conditional probability that team S completes the task given successful coordination. Their product represents the overall success probability, capturing both organizational stability and team-level efficiency. This hierarchical formulation aligns with AEX's reputation system: p_h^{hub} evolves through exponentially weighted updates over historical performance, and $p_{\text{success}}(S)$ is inferred from intra-hub collaboration data.

Collaborative value attribution. Post-task credit is allocated by Shapley value to ensure fair compensation reflecting each agent's marginal contribution (Hua et al., 2025; Yang et al., 2025b):

$$\phi_i = \sum_{S \subseteq A \setminus \{i\}} \frac{|S|!(|A| - |S| - 1)!}{|A|!} \big[v(S \cup \{i\}) - v(S)\big], \tag{8}$$

where $v(S)$ is the coalition value function. Monte Carlo sampling approximates large-scale cases. This ex-post credit distribution ensures fairness within teams and complements the ex-ante efficiency of the auction layer, jointly fostering both truthful bidding and equitable collaboration.

Adaptive mechanism selection. When there are a large number of agents available on the platform, a competitive mechanism can be more appropriate. If there are fewer options or the task has strong requirements for specific agents, a more stable mechanism will be more suitable.

3.4 Market Patterns and Applications

Configurations in Figure 4 correspond to distinct economic patterns observed in real-world markets:

- **Config I (Hub Auction + Direct Assignment)** mirrors competitive service provider markets where platforms compete for projects through bidding (e.g., advertising slot auctions) while maintaining direct control over task execution through predetermined agent assignments. This optimizes platform-level competition while ensuring execution efficiency.
- **Config II (Direct Hub + Agent Auction)** reflects outsourcing platforms where clients directly select service providers who then conduct internal competitive bidding for task execution. This pattern is common in enterprise procurement where trust relationships determine platform selection but competitive dynamics optimize internal resource allocation.
- **Config III (Dual-Level Auction)** handles complex multi-tier markets involving both platform competition and internal resource competition (e.g., supply chain management). This maximizes competitive efficiency at both coordination levels.
- **Config IV (Full Direct Assignment)** represents long-term partnership models where both platform selection and task assignment are predetermined, minimizing transaction costs for recurring collaborations.

This economic correspondence demonstrates AEX's capability to accommodate diverse market structures through adaptive mechanism selection based on market conditions and task characteristics.

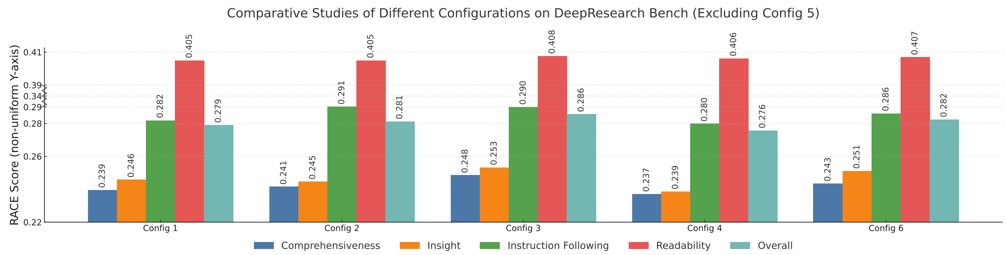

Figure 5: Performance comparison of auction configurations on DeepResearch Bench. Results shown for 5-Arch pool with GPT-4o-mini backbone (OpenAI, 2024a).

4 Empirical Evaluation: Multi-Level Auction Mechanisms

We conduct exploratory experiments to validate AEX's auction mechanisms. Specifically, we investigate three questions: (1) Does hub-level competition improve task outcomes compared to direct assignment? (2) Does adding intra-hub agent competition provide additional gains? (3) How does agent pool composition (architecture diversity and LLM backbone) affect performance?

4.1 Experimental Setup

Benchmark. We use *DeepResearch Bench* (Du et al., 2025), a comprehensive benchmark of PhD-level research tasks, restricting evaluation to 50 English tasks spanning multiple research fields. Evaluation uses **RACE** (Reference-based Adaptive Criteria-driven Evaluation), which reports five dimension scores: *Overall* (aggregate quality), *Comprehensiveness* (coverage breadth), *Insight* (analytical depth), *Instruction Following* (constraint adherence), and *Readability* (presentation clarity). External retrieval is disabled to avoid exogenous variance. Final reports are judged using **gemini-2.5-flash** following the benchmark protocol (Google, 2026).

Configurations. We evaluate 5 configurations (Table 1): (I) hub auction with direct intra-hub assignment; (II) agent auction only; (III) dual-level auction; (IV) direct assignment baseline; (V) capability-first matching.

Implementation. AEX implements: (1) *Task Loader* for structured specifications; (2) *Hub Layer* managing agent pools with bids (p_i, t_i, q_i, c_i); (3) *Auction Layer* scoring bids via $\text{score}(b_i) = \sum_k w_k f_k(b_i)$ with weights w_{cost}=0.3, w_{time}=0.2, w_{quality}=0.3, w_{conf}=0.2; (4) *Execution Layer* for team selection and task execution.

Workflow. At hub level, each hub proposes a team with cost, time, quality, and confidence estimates. AEX selects winners via weighted scoring. The winning hub builds its team by role-skill matching, with staged execution and review cycles.

Agent Pool Construction. We construct agent pools from diverse architectures (Table 2). We evaluate two pool configurations: 5-Arch contains five architectures in Table 2 with GPT-4o-mini as the backbone (OpenAI, 2024a). 7-Arch includes all seven architectures, tested under two backbone settings: (1) GPT-4o-mini for homogeneous evaluation, and (2) Multi-LLM using three LLMs (GPT-4o-mini, GPT-3.5-turbo, GPT-4) to assess LLM heterogeneity effects (OpenAI, 2024c; OpenAI et al., 2024).

Table 2: Agent architecture composition.

Architecture	Pattern	Role
ReAct	Reasoning+Acting	Analysis
Reflection	Self-Critique	Review
Chain-of-Thought	Sequential	Analysis
Planning	Decomposition	Coordinator
Hierarchical	Multi-Level	Manager
Adaptive	Adaptive	Manager
Collaborative	Multi-Perspective	Manager

4.2 Results and Analysis

Table 3, Figure 5 and Figure 6 present results across configurations and pool types.

Q1: Hub-level competition. Configurations with hub-level auctions (Config I, III) consistently outperform those without (Config II, IV). Across all pool types, Config III achieves the highest Overall scores, while Config IV (direct assignment) ranks lowest. This indicates that competitive hub selection improves outcomes.

Q2: Intra-hub agent competition. Comparing Config I (hub auction only) and Config III (dual-level), the additional intra-hub competition in Config III yields marginal gains in homogeneous

Table 3: Performance across agent pool configurations. 5-Arch: 5 architectures with GPT-4o-mini. 7-Arch: all 7 architectures, tested with GPT-4o-mini and Multi-LLM backbones.

Configuration	Pool Type	Comp.	Insight	Inst. Follow.	Read.	Overall
Config I	5-Arch (GPT-4o-mini)	0.2395	0.2458	0.2819	0.4050	0.2790
	7-Arch (GPT-4o-mini)	0.2446	0.2526	0.2869	0.4106	0.2850
	7-Arch (Multi-LLM)	0.2292	0.2279	0.2671	0.3753	0.2612
Config II	5-Arch (GPT-4o-mini)	0.2415	0.2446	0.2913	0.4050	0.2812
	7-Arch (GPT-4o-mini)	0.2344	0.2372	0.2826	0.4008	0.2740
	7-Arch (Multi-LLM)	0.2090	0.2117	0.2525	0.3831	0.2492
Config III	5-Arch (GPT-4o-mini)	0.2485	0.2533	0.2902	0.4076	0.2856
	7-Arch (GPT-4o-mini)	0.2446	0.2493	0.2831	0.4053	0.2816
	7-Arch (Multi-LLM)	0.2360	0.2366	0.2720	0.4040	0.2725
Config IV	5-Arch (GPT-4o-mini)	0.2371	0.2386	0.2798	0.4061	0.2757
	7-Arch (GPT-4o-mini)	0.2333	0.2361	0.2743	0.4025	0.2721
	7-Arch (Multi-LLM)	0.2117	0.2125	0.2531	0.3789	0.2491
Config V	5-Arch (GPT-4o-mini)	0.2434	0.2511	0.2861	0.4069	0.2824
	7-Arch (GPT-4o-mini)	0.2407	0.2471	0.2885	0.4122	0.2822
	7-Arch (Multi-LLM)	0.2353	0.2363	0.2830	0.3821	0.2709

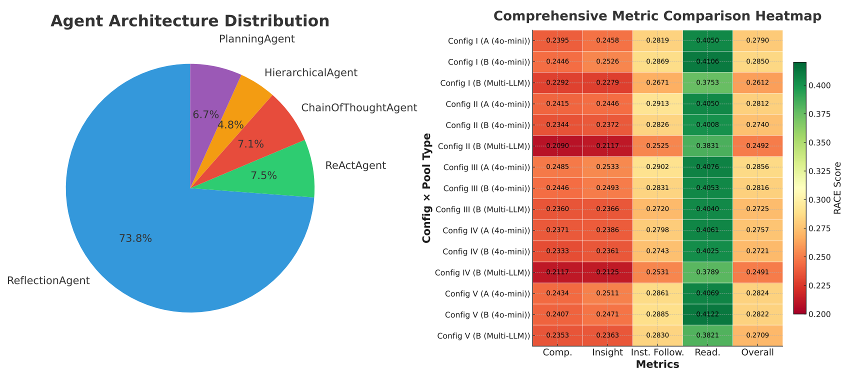

Figure 6: Agent architecture distribution and performance across different agent pool configurations.

pools (GPT-4o-mini) but larger improvements in heterogeneous pools (Multi-LLM). Config II (agent auction only, no hub competition) underperforms Config I, suggesting hub-level selection matters more than agent-level selection alone.

Q3: Pool composition effects. While 7-Arch pools exhibit greater architectural diversity, performance gains are not monotonic with pool complexity. This suggests that more complex agent compositions and model diversity do not automatically yield better outcomes, rather, increased system complexity imposes higher coordination demands among heterogeneous agents, placing greater requirements on the platform's mechanism design to effectively orchestrate collaboration.

These exploratory results provide initial evidence that market-based coordination through AEX improves task outcomes over direct assignment, supporting the "agent-as-actor" paradigm where agents compete through structured auction mechanisms.

5 Conclusion

We present the AEX framework to establish a coherent architectural vision for agent-centric economies, recognizing that realizing such systems requires integrated design across demand specification, capability representation, auction coordination, and value attribution. While our empirical evaluation focuses on validating the core auction mechanisms, the central engine enabling competitive agent selection, we describe the full ecosystem to demonstrate how these mechanisms integrate with complementary infrastructure components. This holistic treatment provides both a validated foundation and a roadmap for comprehensive system development.

REFERENCES

Sunil Adhau, M.L. Mittal, and Abhinav Mittal. A multi-agent system for distributed multi-project scheduling: An auction-based negotiation approach. *Engineering Applications of Artificial Intelligence*, 25(8):1738–1751, 2012. ISSN 0952-1976. doi:https://doi.org/10.1016/j.engappai.2011.12.003. URL https://www.sciencedirect.com/science/article/pii/S0952197611002363.

Saaket Agashe, Yue Fan, Anthony Reyna, and Xin Eric Wang. Llm-coordination: Evaluating and analyzing multi-agent coordination abilities in large language models, 2025. URL https://arxiv.org/abs/2310.03903.

Anthropic. Model context protocol, 2024. URL https://www.anthropic.com/news/model-context-protocol. Accessed: 2025-04-19.

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*, WWW '25, pp. 104–113, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400713316. doi:10.1145/3701716.3715259. URL https://doi.org/10.1145/3701716.3715259.

Junzhe Chen, Xuming Hu, Shuodi Liu, Shiyu Huang, Wei-Wei Tu, Zhaofeng He, and Lijie Wen. LLMArena: Assessing capabilities of large language models in dynamic multi-agent environments. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 13055–13077, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi:10.18653/v1/2024.acl-long.705. URL https://aclanthology.org/2024.acl-long.705/.

Xiang Chen, Yuling Shi, Qizhen Lan, Yuchao Qiu, and Xiaodong Gu. Fed-se: Federated self-evolution for privacy-constrained multi-environment llm agents, 2025. URL https://arxiv.org/abs/2512.08870.

Ye Chen, Pavel Berkhin, Bo Anderson, and Nikhil R. Devanur. Real-time bidding algorithms for performance-based display ad allocation. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '11, pp. 1307–1315, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450308137. doi:10.1145/2020408.2020604. URL https://doi.org/10.1145/2020408.2020604.

Mingxuan Du, Benfeng Xu, Chiwei Zhu, Xiaorui Wang, and Zhendong Mao. Deepresearch bench: A comprehensive benchmark for deep research agents, 2025. URL https://arxiv.org/abs/2506.11763.

Roman J Georgio, Caelum Forder, Suman Deb, Peter Carroll, and Önder Gürcan. The coral protocol: Open infrastructure connecting the internet of agents. *arXiv preprint arXiv:2505.00749*, 2025.

Google. A2a: Agent2agent protocol, 2025. URL https://github.com/google/A2A. Accessed: 2025-04-21.

Google. Gemini models. Google AI for Developers documentation (Gemini API), 2026. URL https://ai.google.dev/gemini-api/docs/models. Includes stable model code gemini-2.5-flash (latest update June 2025); page last updated 2026-01-22 UTC; accessed 2026-02-04.

Sirui Hong, Xiawu Zheng, Jonathan Chen, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, et al. Metagpt: Meta programming for multi-agent collaborative framework. *arXiv preprint arXiv:2308.00352*, 3(4):6, 2023.

Xinyi Hou, Yanjie Zhao, Shenao Wang, and Haoyu Wang. Model context protocol (mcp): Landscape, security threats, and future research directions, 2025. URL https://arxiv.org/abs/2503.23278.

Yun Hua, Haosheng Chen, Shiqin Wang, Wenhao Li, Xiangfeng Wang, and Jun Luo. Shapley-coop: Credit assignment for emergent cooperation in self-interested llm agents, 2025. URL https://arxiv.org/abs/2506.07388.

Gurusha Juneja, Alon Albalak, Wenyue Hua, and William Yang Wang. Magpie: A dataset for multi-agent contextual privacy evaluation, 2025. URL https://arxiv.org/abs/2506.20737.

V. Krishna. *Auction Theory*. Academic Press, 2009. ISBN 9780080922935. URL https://books.google.co.jp/books?id=qW1128ktG1gC.

Jonghan Lim and Ilya Kovalenko. A large language model-enabled control architecture for dynamic resource capability exploration in multi-agent manufacturing systems, 2025. URL https://arxiv.org/abs/2505.22814.

Bang Liu, Xinfeng Li, Jiayi Zhang, Jinlin Wang, Tanjin He, Sirui Hong, Hongzhang Liu, Shaokun Zhang, Kaitao Song, Kunlun Zhu, Yuheng Cheng, Suyuchen Wang, Xiaoqiang Wang, Yuyu Luo, Haibo Jin, Peiyan Zhang, Ollie Liu, Jiaqi Chen, Huan Zhang, Zhaoyang Yu, Haochen Shi, Boyan Li, Dekun Wu, Fengwei Teng, Xiaojun Jia, Jiawei Xu, Jinyu Xiang, Yizhang Lin, Tianming Liu, Tongliang Liu, Yu Su, Huan Sun, Glen Berseth, Jianyun Nie, Ian Foster, Logan Ward, Qingyun Wu, Yu Gu, Mingchen Zhuge, Xiangru Tang, Haohan Wang, Jiaxuan You, Chi Wang, Jian Pei, Qiang Yang, Xiaoliang Qi, and Chenglin Wu. Advances and challenges in foundation agents: From brain-inspired intelligence to evolutionary, collaborative, and safe systems, 2025. URL https://arxiv.org/abs/2504.01990.

Manus. Manus ai platform, 2025. URL https://manus.im/. Accessed: 2025-06-13.

Samuele Marro, Emanuele La Malfa, Jesse Wright, Guohao Li, Nigel Shadbolt, Michael Wooldridge, and Philip Torr. A scalable communication protocol for networks of large language models, 2024. URL https://arxiv.org/abs/2410.11905.

OpenAI. Gpt-4o mini: advancing cost-efficient intelligence. OpenAI (blog post), July 2024a. URL https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/. Accessed: 2026-02-04.

OpenAI. Learning to reason with llms. https://openai.com/index/learning-to-reason-with-llms/, 2024b. Accessed: 2024-04-23.

OpenAI. Gpt-3.5 turbo model. OpenAI API documentation, 2024c. URL https://platform.openai.com/docs/models/gpt-3.5-turbo. Accessed: 2026-02-04.

OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, et al. Gpt-4 technical report, 2024. URL https://arxiv.org/abs/2303.08774.

PwC US. Pwc's ai agent survey 2025. PwC (web report), May 2025. URL https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-agent-survey.html. Accessed: 2025-06-09.

Changle Qu, Sunhao Dai, Xiaochi Wei, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, Jun Xu, and Ji-rong Wen. Tool learning with large language models: a survey. *Frontiers of Computer Science*, 19(8), January 2025. ISSN 2095-2236. doi:10.1007/s11704-024-40678-2. URL http://dx.doi.org/10.1007/s11704-024-40678-2.

Rajesh Ranjan, Shailja Gupta, and Surya Narayan Singh. Loka protocol: A decentralized framework for trustworthy and ethical ai agent ecosystems, 2025. URL https://arxiv.org/abs/2504.10915.

Salesforce. Agentexchange — salesforce, 2025. URL https://www.salesforce.com/agentforce/agentexchange/. Accessed: 2025-06-13.

Ranjan Sapkota, Konstantinos I. Roumeliotis, and Manoj Karkee. Ai agents vs. agentic ai: A conceptual taxonomy, applications and challenges, 2025. URL https://arxiv.org/abs/2505.10468.

Alex Singla, Alexander Sukharevsky, Lareina Yee, Michael Chui, and Bryce Hall. The state of ai in early 2024: Gen ai adoption spikes and starts to generate value. *McKinsey & Company*, May 2024. URL https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024. McKinsey Global Survey on AI.

Peter Stone, Gal A. Kaminka, Sarit Kraus, and Jeffrey S. Rosenschein. Ad hoc autonomous agent teams: collaboration without pre-coordination. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, AAAI'10, pp. 1504–1509. AAAI Press, 2010.

Efosa Udinmwen. Love and hate: Tech pros overwhelmingly like ai agents but view them as a growing security risk. TechRadar (web article), June 2025. URL https://www.techradar.com/computing/artificial-intelligence/love-and-hate-tech-pros-overwhelmingly-like-ai-agents-but-view-them-as-a-growing-security-risk. Accessed: 2025-06-09.

Dennis Ulmer, Elman Mansimov, Kaixiang Lin, Lijia Sun, Xibin Gao, and Yi Zhang. Bootstrapping llm-based task-oriented dialogue agents via self-talk. In *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 9500–9522, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi:10.18653/v1/2024.findings-acl.566. URL https://aclanthology.org/2024.findings-acl.566/.

Jun Wang, Weinan Zhang, and Shuai Yuan. Display advertising with real-time bidding (rtb) and behavioural targeting, 2017. URL https://arxiv.org/abs/1610.03013.

Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Jirong Wen. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6), March 2024. ISSN 2095-2236. doi:10.1007/s11704-024-40231-1. URL http://dx.doi.org/10.1007/s11704-024-40231-1.

Siyuan Wang, Zhuohan Long, Zhihao Fan, Xuanjing Huang, and Zhongyu Wei. Benchmark self-evolving: A multi-agent framework for dynamic LLM evaluation. In Owen Rambow, Leo Wanner, Marianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, and Steven Schockaert (eds.), *Proceedings of the 31st International Conference on Computational Linguistics*, pp. 3310–3328, Abu Dhabi, UAE, January 2025. Association for Computational Linguistics. URL https://aclanthology.org/2025.coling-main.223/.

Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Shaokun Zhang, Erkang Zhu, Beibin Li, Li Jiang, Xiaoyun Zhang, and Chi Wang. Autogen: Enabling next-gen llm applications via multi-agent conversation framework. *arXiv preprint arXiv:2308.08155*, 3(4), 2023a.

Yuxiang Wu, Guanting Dong, and Weiran Xu. Semantic parsing by large language models for intricate updating strategies of zero-shot dialogue state tracking. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 11093–11099, Singapore, December 2023b. Association for Computational Linguistics. doi:10.18653/v1/2023.findings-emnlp.741. URL https://aclanthology.org/2023.findings-emnlp.741/.

Di Xue, Lei Yuan, Zongzhang Zhang, and Yang Yu. Efficient multi-agent communication via shapley message value. In Lud De Raedt (ed.), *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pp. 578–584. International Joint Conferences on Artificial Intelligence Organization, 7 2022. doi:10.24963/ijcai.2022/82. URL https://doi.org/10.24963/ijcai.2022/82. Main Track.

Yingxuan Yang, Qiuying Peng, Jun Wang, Ying Wen, and Weinan Zhang. Llm-based multi-agent systems: Techniques and business perspectives, 2024. URL https://arxiv.org/abs/2411.14033.

Yingxuan Yang, Huacan Chai, Yuanyi Song, Siyuan Qi, Muning Wen, Ning Li, Junwei Liao, Haoyi Hu, Jianghao Lin, Gaowei Chang, Weiwen Liu, Ying Wen, Yong Yu, and Weinan Zhang. A survey of ai agent protocols, 2025a. URL https://arxiv.org/abs/2504.16736.

Yingxuan Yang, Bo Huang, Siyuan Qi, Chao Feng, Haoyi Hu, Yuxuan Zhu, Jinbo Hu, Haoran Zhao, Ziyi He, Xiao Liu, Zongyu Wang, Lin Qiu, Xuezhi Cao, Xunliang Cai, Yong Yu, and Weinan Zhang. Who's the mvp? a game-theoretic evaluation benchmark for modular attribution in llm agents, 2025b. URL https://arxiv.org/abs/2502.00510.

Yingxuan Yang, Mulei Ma, Yuxuan Huang, Huacan Chai, Chenyu Gong, Haoran Geng, Yuanjian Zhou, Ying Wen, Meng Fang, Muhao Chen, Shangding Gu, Ming Jin, Costas Spanos, Yang Yang, Pieter Abbeel, Dawn Song, Weinan Zhang, and Jun Wang. Agentic web: Weaving the next web with ai agents, 2025c. URL https://arxiv.org/abs/2507.21206.

Shuai Yuan, Jun Wang, and Xiaoxue Zhao. Real-time bidding for online advertising: measurement and analysis. In *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising*, ADKDD '13, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450323239. doi:10.1145/2501040.2501980. URL https://doi.org/10.1145/2501040.2501980.

Weinan Zhang, Shuai Yuan, and Jun Wang. Optimal real-time bidding for display advertising. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '14, pp. 1077–1086, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450329569. doi:10.1145/2623330.2623633. URL https://doi.org/10.1145/2623330.2623633.

Ruiwen Zhou, Yingxuan Yang, Muning Wen, Ying Wen, Wenhao Wang, Chunling Xi, Guoqiang Xu, Yong Yu, and Weinan Zhang. Trad: Enhancing llm agents with step-wise thought retrieval and aligned decision, 2024. URL https://arxiv.org/abs/2403.06221.