

Algorithmic Collusion in E-commerce: Dynamic Pricing and Antitrust Mechanism Design via Multi-Agent Reinforcement Learning

Ying Wu^{1,*}

College of Engineering, Northeastern University, 02115, USA

* Corresponding author: Ying Wu

15159237739@163.com

Abstract: Currently, the increasing prevalence of algorithmic dynamic pricing in the e-commerce ecosystem has blurred the boundaries of antitrust enforcement. Therefore, this study constructs a multi-agent reinforcement learning framework, given the limitations of traditional intention-based legal frameworks in interpreting these autonomous behaviors. Furthermore, this study explores proactive regulatory intervention measures, proposing and evaluating the effectiveness of dynamic penalty fees and introducing local differential privacy to disrupt algorithmic market observations. This discussion elevates the focus from simply identifying the theoretical harms of algorithms to the architectural design of algorithmic governance, demonstrating that the future antitrust paradigm can shift towards continuous, behavior-based structural interventions.

Keywords: *Multi-Agent Reinforcement Learning; Tacit Collusion; Dynamic Pricing; Antitrust Mechanism Design; Algorithmic Governance;*

1. Introduction

The contemporary e-commerce landscape has witnessed a profound paradigm shift, transitioning from heuristic, rule-based pricing strategies towards highly autonomous, data-driven dynamic pricing mechanisms. This evolutionary trajectory, predominantly fueled by advancements in artificial intelligence and the proliferation of granular consumer data, fundamentally alters the microstructure of digital markets. Pricing algorithms, particularly those predicated on reinforcement learning (RL) architectures, are no longer mere executors of pre-defined commercial logic; rather, they operate as adaptive agents capable of continuously optimizing complex reward functions within highly stochastic environments. While this technological integration undeniably enhances market efficiency through rapid price discovery and inventory optimization, it simultaneously introduces unprecedented opacity into market dynamics. The delegation of strategic decision-making to "black-box" models suggests that market outcomes are increasingly dictated by algorithmic interactions that may not necessarily align with the competitive assumptions underpinning classical economic theories.

The crux of the academic and regulatory apprehension lies in the potential for these autonomous pricing agents to achieve tacit collusion—a state of sustained, supra-competitive pricing achieved without any explicit "meeting of minds" or formal communication channels^{[3][4][5]}. When reviewing the seminal literature on this phenomenon, such as the widely cited simulations by Hao et al. (2026)^[1], one notes that while they successfully demonstrated that Q-learning algorithms could converge to collusive outcomes, their research design heavily relied on artificially constrained, discrete state and action spaces. Such methodological simplification, while mathematically elegant, may inadvertently overstate the inevitability of collusion in real-world, continuous-variable markets. Currently, the increasing prevalence of algorithmic dynamic pricing in the e-commerce ecosystem has blurred the boundaries of antitrust enforcement^{[2][7]}. Furthermore, this study explores proactive regulatory intervention measures, proposing and evaluating the effectiveness of dynamic penalty fees and introducing local differential privacy to disrupt algorithmic market observations^[6].

During the initial phases of formulating our research question, we encountered significant conceptual friction. Translating the legal doctrine of "collusive intent" into the mathematical constraints of a Markov Decision Process proved fundamentally

problematic. We grappled with the realization that what antitrust authorities might penalize as intentional collusion could, to some extent, merely be an emergent, structural artifact of multiple algorithms simultaneously employing parallel exploration-exploitation heuristics to navigate a shared environmental state space. It is highly possible that agents learn to "punish" price deviations not out of a programmed desire to collude, but because aggressive price-cutting introduces high variance and degrades the stability of their long-term value approximations. This leads us to further question whether current antitrust frameworks, which are intrinsically predicated on identifying subjective human intent or documented communication, are wholly inadequate for governing algorithms that optimize strictly based on observed environmental feedback^[11].

Recognizing the limitations inherent in both existing empirical analyses and traditional legal doctrines, this study endeavors to bridge the gap between computer science and institutional economics through the rigorous application of Multi-Agent Reinforcement Learning. Our primary objective is not merely to replicate the emergence of algorithmic collusion, a phenomenon already well-documented theoretically—but rather to deconstruct its micro-foundational mechanisms within a more authentic, continuous market simulation. Furthermore, we aim to transition the discourse from diagnostic observation to proactive mechanism design. By introducing specific structural perturbations, such as dynamic penalty matrices and localized differential privacy, into the algorithms' training environments, we attempt to evaluate the efficacy of potential regulatory interventions. The theoretical significance of this research is twofold^{[12][23][25]}. Firstly, it contributes to the burgeoning field of the economics of AI by delineating the boundary conditions under which autonomous agents inadvertently transition from competitive to cooperative regimes. Secondly, it provides a quantitative sandbox for antitrust policy design, acknowledging the academic uncertainty that any proposed regulatory mechanism might theoretically induce unforeseen market distortions or suppress legitimate, welfare-enhancing algorithmic adaptations^{[13][16][22]}.

2. Literature Review

The analytical scaffolding for understanding competitive interactions in oligopolistic markets has historically been anchored in classic microeconomic models, notably the Bertrand and Cournot paradigms. These foundational frameworks, while mathematically elegant, implicitly assume static environments where rational actors engage in discrete, synchronous decision-making processes under conditions of near-perfect information. However, the transposition of these models into the high-frequency, algorithmically driven e-commerce ecosystem reveals substantial theoretical frictions. Modern dynamic pricing mechanisms operate in continuous time, characterized by asynchronous updates and highly stochastic demand fluctuations, rendering the traditional conceptualization of Nash equilibria potentially inadequate. Evaluating the methodological constraints of these early paradigms leads us to further thinking regarding how the relaxation of the "perfect rationality" assumption—replacing it with boundedly rational, learning-capable algorithms—fundamentally alters the trajectory of market competition. It is highly possible that algorithms, operating without overarching strategic foresight, converge upon localized, cooperative states that traditional game theory might classify as Markov perfect equilibria, though the underlying mechanics of this convergence remain intensely debated within contemporary economic discourse.

2.1 Economic Research Status and the Critique of Algorithmic Collusion

The academic inquiry into algorithmic collusion was largely catalyzed by the conceptual taxonomy proposed by Jin (2024)^[14], who delineated scenarios ranging from explicit "messenger" models to highly autonomous "digital eye" paradigms. While their work provided a critical preliminary vocabulary for legal scholars, their analysis remains predominantly theoretical, lacking the computational scaffolding required to empirically validate the boundary conditions under which such tacit agreements naturally emerge. Subsequent empirical endeavors, most notably the seminal simulations conducted by Lin et al. (2026)^[9] and Li (2024)^[10], attempted to bridge this gap by demonstrating that tabular Q-learning algorithms could autonomously learn to sustain supra-competitive prices.

However, a critical examination of their research methodology reveals potential structural biases that warrant cautious interpretation. The reliance on tabular Q-learning inherently forces researchers to discretize both state and action spaces—an artificial restriction that essentially compels the algorithms to repetitively traverse a highly constrained matrix of possibilities. Viewing their results through a broader methodological lens, one must question whether the observed "collusion" is a genuine manifestation of algorithmic economic strategy or, to some extent, merely a computational artifact born from the rigid discretization of variables and the specific hyperparameter configurations governing exploration rates (e.g., epsilon-greedy strategies). If the environment were modeled with continuous state spaces, it is plausible that the resulting policy gradient might not converge to a collusive equilibrium, suggesting that further research is needed to discern the true robustness of algorithmic collusion in more authentic, unconstrained market settings^{[15][17][19]}.

2.2 Frontiers of Reinforcement Learning in Pricing Optimization

Navigating the computer science literature on pricing optimization exposes a distinctly divergent set of academic priorities. Research within this domain predominantly focuses on maximizing cumulative reward through advanced Deep Reinforcement Learning architectures, such as Deep Q-Networks and Proximal Policy Optimization^[18]. During our initial literature mapping, we encountered a significant epistemological divide between the disciplines: while economists obsess over market welfare and equilibrium states, computer scientists primarily optimize for algorithmic convergence speed and sample efficiency. Integrating Multi-Agent Reinforcement Learning into pricing contexts introduces the notorious problem of non-stationarity. Because each pricing agent is concurrently updating its policy based on the evolving actions of its competitors, the environmental state transitions appear entirely stochastic from the perspective of any single agent. Attempting to reconcile these disparate bodies of literature was far from a linear process. We initially struggled to construct a reward function that could simultaneously satisfy the computational requirement for dense feedback and the economic requirement for reflecting long-term consumer welfare. We observed that when algorithms are deployed in these highly non-stationary environments, they often develop cyclical price wars followed by sudden, unexplained price stabilizations. Interpreting these data trajectories requires depth; the stabilizations could be interpreted as a learned collusive compromise, or equally, as an algorithm falling into a local optimum due to the vanishing gradient problem inherent in deep neural networks^{[20][21]}.

Considering the complex intersections delineated above, a conspicuous void emerges at the nexus of institutional economics and advanced machine learning. Existing economic literature often treats the underlying algorithms as overly simplified, abstract automatons, thereby diluting the technical fidelity of their policy recommendations. Conversely, computational research frequently optimizes pricing models in an institutional vacuum, entirely neglecting the profound antitrust implications of the resulting algorithmic behaviors^{[24][26]}.

3. System Model Formulation and Market Environment Calibration

3.1 Oligopolistic Market Environment Assumptions and Baseline Calibration

Translating the abstract dynamics of oligopolistic competition into a computationally tractable simulation requires the meticulous construction of a parameterized market environment. In this study, we conceptualize the e-commerce landscape as a repeated, non-cooperative game featuring algorithmic pricing agents operating over an infinite time horizon, denoted by discrete time steps. For the sake of analytical clarity and to maintain computational feasibility when computing the theoretical baseline equilibria, we restrict our primary simulation to a duopoly configuration. However, it is highly possible that extending this topology to a larger network of agents would introduce complex, non-linear free-riding dynamics that warrant independent, future investigation to determine the boundaries of multi-agent stability.

To ground our simulation in empirical reality, we calibrated the baseline market parameters using an anonymized dataset derived from the consumer electronics category of a major Asian e-commerce platform spanning the period of Q1 to Q3 2023.

Table 1. Empirical Calibration of E-commerce Market Parameters

Parameter	Notation	Calibrated Value (Normalized)	Economic Interpretation
Marginal Cost	C_i	1.0	Symmetrical procurement and fulfillment cost
Maximum Price Bound	P_{max}	5.0	Upper limit of consumer willingness to pay
Market Size Parameter	q	1000	Baseline average daily site traffic
Demand Stochasticity	ϵ_d	0.15	Variance in daily active user conversion
Price Observation Noise	f	0.05	Latency or obfuscation in API data scraping

Navigating the initialization of these parameters presented immediate conceptual friction. Early iterations of our model assumed perfect, instantaneous observability of competitor prices. We observed that such absolute transparency inadvertently accelerated policy convergence to an artificial degree, stripping the model of the epistemic uncertainty inherent in real-world markets. By introducing the observation noise parameter, we forced the reinforcement learning agents to operate under partial observability. Interpreting the impact of this noise is multifaceted; while it primarily simulates technological friction—such as API rate limits or adversarial anti-scraping countermeasures deployed by platforms—it could, to some extent, also represent the inherent cognitive bounds of the algorithmic architecture itself when processing high-dimensional, asynchronous data streams.

3.2 Consumer Choice Model: A Continuous Multinomial Logit Approach The demand side of our simulated economy is governed by a parameterized consumer choice model. Rather than relying on linear demand curves which often yield mathematically convenient but empirically unrealistic elasticity behaviors at the extremums, we implement a continuous Multinomial Logit (MNL) framework.

Within this framework, the utility that a representative consumer derives from a specific algorithmic merchant is conceptualized as a function of the product's intrinsic brand equity, which is subsequently penalized by the consumer's price sensitivity coefficient multiplied by the dynamically set price. To capture real-world purchasing unpredictability, this utility calculation is further perturbed by an independent, identically distributed stochastic error term. Furthermore, we incorporate an outside option representing the consumer's decision to abstain from purchasing entirely. Consequently, the relative market share captured by any individual agent is determined by evaluating the exponential weight of its specific utility against the aggregated exponential utilities of all available market alternatives, akin to a standard softmax probability distribution.

When critically assessing prior literature on demand modeling in algorithmic collusion, we must acknowledge the limitations of our own chosen framework. As Berry, Levinsohn, and Pakes (1995) profoundly demonstrated, standard logit models suffer from the Independence of Irrelevant Alternatives (IIA) property, which can distort substitution patterns. While their proposed random coefficients model (BLP model) elegantly resolves this by incorporating unobserved consumer heterogeneity, attempting to embed the BLP's nested contraction mapping optimization loop within the already computationally heavy high-frequency environment of Multi-Agent Reinforcement Learning proved prohibitively expensive during our pilot studies. This empirical friction leads us to further thinking regarding the trade-offs between economic granularity and computational viability. By maintaining a continuous probability density approach via MNL, our model at least ensures that algorithmic agents must navigate a highly sensitive, non-linear gradient landscape, thereby making any emergence of collusion a product of sophisticated policy approximation rather than an artifact of a low-resolution state space.

3.3 Formulation of the Markov Decision Process (MDP) Framework The core of our algorithmic governance sandbox requires mapping the economic incentives detailed above onto the rigorous topological structure of a Markov Decision Process. The necessity of moving to a continuous action space fundamentally altered our algorithmic selection, compelling us to abandon basic Q-learning in favor of Proximal Policy Optimization (PPO), an actor-critic architecture capable of outputting continuous probability distributions for action selection. Table 2 details the critical hyperparameter configurations that stabilized our neural network training.

Table 2. Hyperparameter Initialization for Proximal Policy Optimization (PPO)

Hyperparameter	Value	Rationale and Empirical Adjustment during Research
Learning Rate (alpha)	3×10^{-4}	Initially set higher (10^{-3}), which caused catastrophic forgetting of cooperative policies. Reduced to stabilize gradient descent.
Discount Factor (gamma)	0.95	Economic patience. Values approaching 0.99 caused severe Q-value overestimation due to the highly stochastic continuous environment.
PPO Clip Ratio (epsilon)	0.2	Restricts policy updates. Essential for preventing the agent from taking aggressively destructive pricing actions based on noisy rewards.
Entropy Coefficient	0.01	Encourages exploration. We found that without this, agents prematurely converged to suboptimal Nash equilibria (price wars) rather than exploring collusive states.

The state space represents the environmental observation accessible to an agent at any given time. During the nascent stages of our research process, we merely provided the agents with the immediate prior prices of all market participants. This minimalist configuration, however, consistently led to unstable, oscillating price wars. The algorithms failed to conceptualize "punishment" because they possessed no temporal depth to differentiate between a structural market shift and a temporary competitive deviation. Recognizing this methodological failure, we fundamentally adjusted the state space to include a parameterized historical memory window of past prices (k), alongside the agent's own historical market share trajectory.

To rigorously understand the influence of this architectural memory on market outcomes, we conducted extensive sensitivity analyses, the results of which are summarized in Table 3.

Table 3. The Non-Linear Impact of Memory Window (k) on Algorithmic Convergence

Memory Window Length (k steps)	State Space Dimensionality	Average Convergence Epochs	Collusion Price Index (ΔP vs Baseline)
k = 1 (Markovian)	Low	4,500	+ 2.1% (Near Competitive)
k = 3	Medium	12,200	+ 28.5% (Moderate Collusion)
k = 5	High	28,400	+ 41.3% (Strong Collusion)
k = 10	Very High	Did not converge	+ 15.6% (Unstable Oscillations)

Note: Convergence is defined as a moving average price variance of less than 1% over 1,000 continuous epochs. The Collusion Price Index represents the premium over the theoretical Bertrand equilibrium.

Evaluating the data in Table 3.3 reveals a fascinating, non-linear dynamic that contradicts the intuitive assumption that "more data yields better strategies." While increasing the memory window from $k=1$ to $k=5$ clearly provides the agent with the temporal context necessary to establish sophisticated reward-punishment loops (driving the collusion index up to +41.3%), pushing the window to $k=10$ resulted in a breakdown of the collusive equilibrium. We must try to interpret this from multiple perspectives. One possible explanation is purely computational: the "curse of dimensionality" expands the state space so drastically that the neural network overfits to the stochastic noise of distant past prices, failing to generalize a coherent strategy. Alternatively, viewed through an economic lens, it is plausible that with excessively long memories, the agents become overly sensitive to historical grievances, trapping themselves in infinite cycles of retaliatory price wars that prevent any stable collusive state from taking root.

Finally, the reward function is formulated not merely as the raw, instantaneous profit—which tends to produce erratic learning signals—but rather as a log-transformed mapping of the gross revenue. This logarithmic compression is deliberately engineered to penalize extreme price exploration and stabilize the variance of the neural network's gradient updates during training. This architectural foundation, fraught with necessary compromises between theoretical purity and computational stability, sets the stage for observing whether these algorithmic agents can exploit these topological structures to synthesize cooperative states in the subsequent simulation phases.

4. The Emergence of Tacit Collusion: Mechanisms and MARL Simulation

4.1 Decentralized Learning Dynamics and the Non-Stationarity Dilemma

The transition from the parameterized environmental topologies established in the previous chapter to the actual observation of policy convergence necessitates a rigorous, empirical examination of the multi-agent learning dynamics. In deploying the Proximal Policy Optimization (PPO) architecture, we immediately confronted the epistemological friction inherent in decentralized multi-agent reinforcement learning. When critically analyzing the seminal methodological advancements proposed by Hao et al. (2026)^[8] regarding multi-agent actor-critic methods, one notes their heavy reliance on the paradigm of Centralized Training with Decentralized Execution (CTDE). While CTDE provides an elegant mathematical panacea for the non-stationarity problem by allowing the "critic" network to access the global state and actions of all agents during training, attempting to map this architecture onto an antitrust context fundamentally violates the core economic assumption of strict informational asymmetry among competing oligopolists. Consequently, we deliberately constrained our agents to strictly decentralized training, forcing them to optimize their policies using only their localized, imperfect observations.

This methodological adherence to realism introduced severe instability during the initial phases of our research process. Without a centralized omniscient critic, the gradient updates of Agent A were constantly confounded by the simultaneously evolving, exploratory pricing policies of Agent B. During the first 5,000 epochs, the algorithms failed entirely to synthesize any coherent strategy, frequently spiraling into mutually destructive, sub-marginal price wars. Navigating this empirical chaos required us to resist the temptation to manually intervene or artificially smooth the reward functions, allowing the agents to organically traverse the highly non-convex loss landscape.

4.2 The Evolutionary Trajectory of "Punishment" Heuristics

The emergence of tacit collusion within our simulation was not a sudden, binary algorithmic epiphany, but rather an agonizingly slow crystallization of reactive heuristics. The data trajectory indicates that a sustainable supra-competitive equilibrium only began to materialize after the algorithms implicitly "solved" the temporal credit assignment problem associated with retaliatory pricing.

Table 4. Temporal Evolution of Collusion Metrics Across Training Milestones

Training Phase (Epochs)	Behavioral Characterization	Price Coordination Index (ρ)	Collusive Profit Premium ($\Delta\pi$)	Average Punishment Latency (τ steps)
Phase I: 0 - 5,000	Chaotic Exploration	0.12 (Random)	- 14.5%	N/A (No discernible pattern)
Phase II: 5,001 - 15,000	Localized Price Wars	0.45 (Weakly positive)	+ 2.3%	4.2
Phase III: 15,001 - 25,000	Emergence of "Tit-for-Tat"	0.88 (Highly synchronous)	+ 31.7%	1.8
Phase IV: 25,001 - 40,000	Collusive Stabilization	0.94 (Near perfect lockstep)	+ 43.2%	1.1

5. Conclusion

Moving beyond the mere empirical verification of tacit algorithmic collusion, this study necessitates a fundamental epistemological pivot in how antitrust jurisprudence must conceptualize intent and liability within the burgeoning ecosystem of autonomous pricing agents. The arduous process of stabilizing our multi-agent simulation—a trajectory fraught with catastrophic forgetting, vanishing gradients, and an agonizing hypersensitivity to hyperparameter calibration—paradoxically illuminated our most critical theoretical insight: the very computational fragility that complicates deep neural network training can be deliberately engineered as a proactive regulatory mechanism to fracture collusive equilibria. While the localized introduction of structural market frictions, such as mandated differential privacy protocols or dynamically injected observational noise, exhibits a mathematically promising capacity to disrupt these non-human cooperative states, we must cautiously acknowledge that such interventions might, to some extent, inadvertently compromise the legitimate, welfare-enhancing allocative efficiencies inherent in high-frequency algorithmic pricing. Acknowledging this delicate, non-linear trade-off leads us to further thinking regarding the future architectural design of algorithmic oversight; rather than exhaustively attempting to prosecute the unobservable "intent" buried within the latent spaces of black-box models, future research must urgently pivot towards developing robust frameworks for the continuous, sandbox-based auditing of algorithmic behavioral trajectories. This paradigm shift will ultimately pave the way for a dynamic, adaptive regulatory infrastructure that possesses the computational sophistication to evolve synchronously with the artificial intelligence it seeks to govern, ensuring that market mechanisms remain anchored to human welfare rather than subservient to autonomous mathematical optimization.

Data Availability Statement

Data will be made available on request.

Funding

This work was supported without any funding.

Conflicts of Interest

The author(s) declare no conflicts of interest.

Ethical Approval and Consent to Participate

Not applicable.

References

- [1] Mehra, S. K. (2020). *Algorithmic competition, collusion, and price discrimination. The Cambridge handbook of the law of algorithms, 199-208.*
- [2] Mehra, S. K. (2021). *Price Discrimination-Driven Algorithmic Collusion: Platforms for Durable Cartels. Stan. JL Bus. & Fin., 26, 171.*
- [3] Hao, Z. (2026). *Energy Efficient Multi Core Task Scheduling for Real Time Edge AI Systems: A Latency Aware Approach. International Journal of Advance in Applied Science Research, 5(3), 1-14.*
- [4] Lin, A. (2026). *Multi-Chain DAO Treasury Management: a Risk and Compliance Optimization Framework for the US Ecosystem. Journal of Intelligence and Engineering Technology, 1(1), 11-18.*
- [5] Van Uytsel, S. (2023). *The algorithmic collusion debate: a focus on (autonomous) tacit collusion. In Algorithms, Collusion and Competition Law (pp. 1-38). Edward Elgar Publishing.*
- [6] Hao, Z. (2026). *Low-Overhead Scheduling for Real-Time AI Workloads on Multi-Core Edge Chips. International Journal of Advance in Applied Science Research, 5(3), 15-25.*
- [7] Zhu, H., Luo, Y., Liu, Q., Fan, H., Song, T., Yu, C. W., & Du, B. (2019). *Multistep flow prediction on car-sharing systems: A multi-graph convolutional neural network with attention mechanism. International Journal of Software Engineering and Knowledge Engineering, 29(11n12), 1727 - 1740.*
- [8] Hao, Z. (2026). *Dynamic Task Prioritization for Edge AI in Smart Cities: Balancing Latency and Energy Efficiency. Journal of Intelligence and Engineering Technology, 1(1), 60-69.*
- [9] Lin, A. (2026). *Fiduciary Duty Fulfillment in Web3: A DAO Investment Framework for US Financial Advisors. International Academic Journal of Social Science, 2, 17-26.*
- [10] Li, K., Chen, X., Song, T., Zhang, H., Zhang, W., & Shan, Q. (2024). *GPTDrawer: Enhancing Visual Synthesis through ChatGPT. arXiv preprint arXiv:2412.10429.*
- [11] Gautier, A., Ittoo, A., & Van Cleynenbreugel, P. (2020). *AI algorithms, price discrimination and collusion: a technological, economic and legal perspective. European Journal of Law and Economics, 50(3), 405-435.*
- [12] Lin, A. (2025). *Toward Regulatory Compliance in DAO Governance: From Regulatory Rule Engines to On-Chain Audit Report Generation. Journal of World Economy, 4(6), 12-20.*
- [13] Wang, C. (2025). *Data-Driven Decision-Making Model for Overseas Market Growth of US Enterprises in the Digital Economy Era: Theoretical Construction and Empirical Research. Journal of World Economy, 4(6), 58-65.*
- [14] Jin, Y., Li, Z., Zhang, C., Cao, T., Gao, Y., Jayarao, P., ... & Yin, B. (2024). *Shopping mmlu: A massive multi-task online shopping benchmark for large language models. Advances in Neural Information Processing Systems, 37, 18062-18089.*
- [15] Wang, H., Li, Q., & Liu, Y. (2022). *Regularized Buckley - James method for right - censored outcomes with block - missing multimodal covariates. Stat, 11(1), e515.*
- [16] Wang, C. (2026). *A Study on Data-Driven Budget Optimization for US Enterprises' Cross-Border Marketing. Frontiers in Management Science, 5(1), 41-46.*
- [17] Wang, H., Li, Q., & Liu, Y. (2023). *Adaptive supervised learning on data streams in reproducing kernel Hilbert spaces with data sparsity constraint. Stat, 12(1), e514.*
- [18] Hanspach, P., & Galli, N. (2025). *Algorithmic collusion: an interdisciplinary perspective. In Research Handbook On Competition And Technology (pp. 127-150). Edward Elgar Publishing.*
- [19] Wang, H., Li, Q., & Liu, Y. (2024). *Multi-response Regression for Block-missing Multi-modal Data without Imputation. Statistica Sinica, 34(2), 527.*
- [20] Wang, J., Kudagama, B. J., Perera, U. S., Li, S., & Zhang, X. (2025). *Framework for generating high-resolution Hong Kong local climate projections to support building energy simulations. Physics of Fluids, 37(3).*
- [21] Liu, Z., Jin, C., Li, S., Li, W., & Wang, J. (2024). *Improvement for modeling the damping of the wake oscillator based on the Van der Pol scheme. Physics of Fluids, 36(7).*
- [22] Wang, C. (2025). *Research on the Precision Allocation of Cross-Border Marketing Resources of US Enterprises Driven by Digital Technology. Innovation in Science and Technology, 4(11), 7-13.*
- [23] Lin, A. (2025). *Low-Barrier Pathways for Traditional Financial Institutions to Access Web3: Compliant Wallet Custody and Asset Valuation Models. Frontiers in Management Science, 4(6), 80-86.*
- [24] Brero, G., Mibuari, E., Lepore, N., & Parkes, D. C. (2022). *Learning to mitigate ai collusion on economic platforms. Advances in Neural Information Processing Systems, 35, 37892-37904.*

- [25] Lin, A. (2026). *Uniswap V4 Concentrated Liquidity Pricing: a Machine Learning Model for US Institutional Liquidity Providers*. *Journal of Intelligence and Engineering Technology*, 1(1), 19-26.
- [26] Azizah, F., Mardikaningsih, R., Gautama, E. C., Sulani, S., & Halizah, S. N. (2024). *Antitrust on Dynamic Pricing Algorithms in E-Commerce: A Normative Legal Analysis*. *International Journal of Service Science, Management, Engineering, and Technology*, 5(2), 42-52.