Applied Computer Science 2026, vol. 22, no. 2, pp. 126–137
https://doi.org/10.35784/acs_9272

Submitted: 2026-02-27 | Revised: 2026-04-29 | Accepted: 2026-05-05

Keywords: autonomous agents, algorithmic pricing, digital markets, competition, AI regulation

Elmira KYDYRBAYEVA [1], *Balhiya SHOMSHEKOVA* [1*], *Asset ABZHAKOV* [2], *Ainur ASHIMOVA* [3], *Assel NURTAYEVA* [1]

[1] Zhetysu University named after I. Zhansugurov, Kazakhstan, elmirocka541@gmail.com, shomshekovab@outlook.com, ak.nurtaeva@gmail.com
[2] Academy of Public Administration under the President of the Republic of Kazakhstan, Kazakhstan, asset.abzhakov@apa.kz
[3] Almaty Management University, Kazakhstan, a.ashimova@almau.edu.kz
[*] Corresponding author: shomshekovab@outlook.com

Autonomous AI agents in digital markets: Economic implications for competition, pricing, and regulation

Abstract

The competitive process for price-setting in digital markets is being dramatically altered through the use of autonomous artificial intelligence (AI) to make pricing decisions on behalf of humans. These systems operate autonomously and interact with each other in continuous cycles. They react in real time to market data and adapt their pricing strategies accordingly. This research analyzes the effects of various combinations of market transparency and algorithmic autonomy on price behavior, the competitive process, and consumer welfare outcomes. The analysis is conducted using a controlled simulation model that compares four pricing regimes: human-supervised, fully autonomous, isolated, and mediated. The results show that the degree of market transparency and the degree of platform oversight of AI decision-making have a far greater impact on the market's final outcome than the level of algorithmic autonomy. Some configurations increase efficiency and profit while increasing the risk of coordination failure, volatility, and concentration. In addition, the findings demonstrate inherent structural trade-offs between market efficiency, market stability, and competitive processes in markets where AI is used as an autonomous decision-making agent. These results also highlight the limitations of attempting to regulate AI through intent-based regulations and provide insights into how autonomous AI decision-making agents alter the structure of digital markets.

1. INTRODUCTION

Autonomous artificial intelligence systems increasingly dominate digital markets by making pricing decisions, as autonomous agents use algorithms to continuously adjust prices based on supply-and-demand signals and interact with other automated pricing agents on online platforms to complete transactions (Calvano et al., 2020). In other words, algorithm-based pricing systems are shifting market outcomes from human-driven to algorithm-driven pricing (Goldfarb & Tucker, 2019) through the continued participation of algorithm-based pricing agents in competitively pricing their products in digital markets.

Autonomous agents' algorithmic pricing presents many challenges, both theoretical and practical, particularly because autonomous algorithms can independently learn pricing strategies and potentially engage in coordinated behavior without explicit human intent (Harrington, 2018). As an example of challenges posed by using algorithmic pricing, researchers demonstrate how automated pricing creates situations resulting in higher price volatility for buyers, creates opportunities to communicate and collaborate with one another without violating competition laws, and has dramatic changes in competitive market dynamics even if you do not have any direct communication or coordination with one another (Klein, 2021). Meanwhile, the expansive use of high-frequency feedback, combined with greater transparency, will further compound their impact. In addition, the current research either examines a single pricing mechanism or a specific design, thereby limiting our ability to gain a comprehensive understanding of the extent to which different degrees of autonomy and levels of control affect overall market outcomes.

Digital platforms do not just act as facilitators; they actively shape the overall structure of pricing on the platform through information disclosure, algorithmic constraints, and centralized control (Agrawal et al., 2019). While creating a stable, certain price outcome, digital platforms can also add new risk to the marketplace through coordination and have different distributional effects. Research on the economic impact of platform-mediated governance in markets with a predominance of autonomous pricing agents has not been conducted to date.

The intent of this study is to empirically explore and develop an understanding of the economic implications of autonomous AI agents in digital marketplaces, based on the architecture (platform-mediated versus human-supervised) within which the agents conduct their pricing activities. This comprehensive study is designed as a supervised simulation, comparing the pricing policies of sellers using human-assisted pricing models with those of sellers setting prices autonomously, as well as all autonomous sellers in isolation. This allows us to analyze how pricing dynamics evolve over time, the implications for seller profits, overall price volatility, and the overall impact on consumer welfare (Bajari et al., 2019).

The contribution of this research is to provide a systematic empirical analysis of both the autonomy of the autonomous AI pricing agents and the various transparency levels that exist in the marketplace by utilizing the proposed experimental framework in all the evaluations. The results will provide substantive new insights into the trade-off between efficiency, stability, and competition in an AI-driven digital marketplace. The study demonstrates the relevance of this research for pricing algorithm design, platform governance, and regulatory policy in digital economies.

2. LITERATURE REVIEW

Spulbar (2025) investigated the pricing strategies employed by rule-based agents that operate in a competitive online marketplace. Findings revealed that simple methods of price-matching and undercutting could respond mechanically to competitors' prices; however, these strategies are unable to adapt to fluctuating demand conditions because they lack learning mechanisms. As such, rule-based methods perform poorly in dynamic or transient environments, underscoring the need for adaptive pricing frameworks.

Kastius and Schlosser (2022) used a reinforcement learning approach to solve dynamic pricing problems, employing Q-learning in competitive environments. The results yielded improved profit optimization compared with static pricing rules; however, convergence rates were slower when multiple agents were simultaneously adapting their strategies. Additionally, sensitivity to exploration parameters remained an issue when utilizing RL-based methods. Hence, these results demonstrate the ongoing scalability problems facing early learning-based pricing approaches.

Further support for the need for adaptive pricing methods is provided by Groeneveld et al. (2024), who examine self-learning agents operating in competitive recommerce markets. Their findings indicate that stable learning trajectories can emerge when model parameters are properly calibrated, and agents repeatedly interact in structured market environments. However, when demand signals are noisy or informational conditions become less reliable, adaptive agents exhibit weaker policy performance and greater instability in pricing outcomes. These results suggest that the development of effective autonomous pricing strategies requires careful methodological calibration and continued refinement of learning architectures.

Additionally, deep reinforcement learning has been introduced into the study of pricing to address complex, high-dimensional decision-making environments (Yavuz & Kaya, 2024). Results from this research indicate that deep reinforcement learning models can process richer market information and autonomously learn adaptive pricing strategies. However, the resulting pricing policies become significantly more difficult to interpret and regulate as they evolve, creating challenges for transparency and regulatory oversight. This highlights the need for comparative research between constrained and unconstrained autonomous agents operating under different governance conditions.

Applying policy gradient approaches to continuous pricing allowed Shahzad et al. (2025) to examine the implications for time frames associated with demand shifts. This revealed that agents using policy gradient approaches will adapt to changing demand significantly faster than those using other forms of agent-based pricing strategies, but will also experience larger price fluctuations due to agent exploration. The remaining issue of trade-offs between increased adaptability and stability continues to underscore the effectiveness of establishing governance provisions for the automated pricing mechanisms employed by conditional-demand agents.

Studies using algorithmic collusion simulations (C. Li & Virtosu, 2025) showed that reinforcement learning agents converged on high-price regimes without engaging in coordinated pricing. Therefore, the absence of an intentional attempt to collude poses a challenge to existing views on antitrust regulation (Q. Li, 2025). These findings emphasized the need for revised regulatory perspectives that consider collective or collusive pricing regimes that may develop from the dynamic behaviors of conditional-demand agents utilizing learning-based methodologies.

Gerpott and Berends (2022) explored how transparency affects pricing dynamics by increasing the speed of learning and decreasing uncertainty through full observations while also increasing the potential for coordinated behavior among agents. Conversely, limited observability will reduce coordination yet increase volatility in the overall market. Transparency is a significant design parameter for price dynamics.

Krarup and Horst (2023) focused on algorithmic pricing and highlighted how delayed and/or noisy feedback from suppliers leads to slower convergence and unstable price trajectories as agents work to infer their real demand signals. The relationship between delay and learning dynamics has yet to be studied in more detail; therefore, more robust evaluations are needed.

Wu (2023) introduced platform-mediated pricing mechanisms as a new market governance model that enables platforms to govern markets through recommended prices and constraints that mitigate extreme outcomes and create correlated activity among the parties involved in transactions. As a result, the level of competition between agents may be diminished. The platform's role as a market designer will become clearer when the implications of this type of activity have been documented.

Normann and Sternberg (2023) investigate pricing behavior in hybrid laboratory markets where human decision-makers interact with algorithmic pricing agents. Their findings indicate that the presence of algorithmic delegation significantly alters competitive dynamics and price levels compared to purely human markets. While structured constraints and delegation mechanisms may influence coordination patterns, algorithmic participation also changes profit distribution and adaptive responses within the market. These results highlight the structural trade-off between control and efficiency in hybrid pricing environments and underscore the need for systematic comparison of alternative governance regimes.

Lancieri et al. (2025) examined the welfare implications of algorithmic pricing. The results show that, under adaptive strategies, sellers achieve higher profits while consumers experience a decrease in consumer surplus at higher levels of coordination. Thus, the distributional effects are significant. Therefore, the welfare aspects of pricing must be analyzed beyond merely comparing prices at each point of sale with the average.

Schauer and Schnurr (2023) evaluated the robustness and sensitivity of multi-agent pricing systems. Results vary greatly depending on an agent's learning rate and exploration policy. Even minor changes to market parameters could cause the market to experience entirely different regimes; thus, a single-run study may not be generalizable. Robust evaluations should be used to form valid conclusions.

Lastly, Comunale and Manera (2024) discussed the regulatory challenges presented by autonomous pricing agents. Traditional competition-regulated law does not work well because it fails to consider opacity or attribution when reviewing anticompetitive behavior. Various proposals to address this issue have been offered, the success of which relies heavily on the design and information structure of the agent in question. Therefore, integrated economic and regulatory analyses are needed to adequately evaluate these challenges.

Existing research shows that autonomous pricing agents can improve market efficiency, stability, and competitive dynamics. However, most studies focus on isolated aspects such as reinforcement learning performance, algorithmic coordination, or regulatory implications. Limited attention has been given to how the interaction between algorithmic autonomy and market transparency jointly influences price volatility, coordination risks, and welfare outcomes. Understanding these combined effects is essential for identifying pricing regimes that balance efficiency, stability, and competition. To address this gap, this study conducts a controlled simulation-based comparative analysis of different levels of algorithmic autonomy and market transparency, examining their impact on pricing dynamics, competition, and welfare outcomes.

3. MATERIALS AND METHODS

Digital markets behave as repeated pricing games in which multiple competing sellers offer perfectly substitutable products. Each seller is modeled as an independent agent making price decisions; sellers update their price levels after each round of competition to reflect actual demand for their product, as well as the profits they made during that round and earlier rounds of competition (based on historical data). The main driver of market demand is price, and market demand is sensitive to relative price levels. In addition to market demand, there is a stochastic component in the data that introduces uncertainty and inaccuracy into estimates of actual demand ("noise"). While marginal costs are assumed to be constant, profit is defined as total revenue less total production costs. Each pricing simulation is run for a maximum of 1,000 rounds to capture both short- and long-term behavior; results from multiple simulations are averaged to obtain a representative dataset.

Four experimental conditions are included: human-supervised independent pricing agents with bounded price changes; fully autonomous pricing agents; independent pricing agents who have limited knowledge about their competitors; and pricing agents working under platform-mediated control. Pricing decisions will be modified according to adaptive learning rules based on the amount of operational price uncertainty (i.e., recent history and current round of activities). The goal of each agent is to maximize profit over the long run. Pricing outcomes are evaluated based on average prices, price dispersion, price volatility, seller profit, and consumer surplus. To evaluate performance and identify gaps in the data distribution, distributional and convergence statistics are used. Key variables associated with learning and information are used to conduct sensitivity analyses.

To establish computational reproducibility and provide an overview of how the adaptive learning rules worked in practice, the following section outlines the code used to implement the AI agents. No proprietary AI APIs or existing reinforcement learning libraries (like Ray RLlib) were used for several reasons; for instance, proprietary AI APIs often act as black boxes that do not offer insight into the underlying economics, making it difficult to understand what is happening under the hood. In its place, a custom multi-agent environment was programmed from scratch using Python version 3.9.7. In addition, all computations were done using the NumPy (v1.21.2) and SciPy (v1.7.1) packages.

For the autonomous agents, the approach taken was the use of Q-learning with a table that followed the theory framework laid out by Calvano et al. (2020) for repeated algorithmic pricing games. At each step t, agent i observes the state s_t, which is defined by the discrete prices chosen by all competing agents in the previous round, and selects a new price $p_{i,t}$ from a finite grid. The Q-values corresponding to the expected cumulative profits for actions in states get updated using the following Bellman equation (1):

$$Q_i(s_t, p_{i,t}) \leftarrow (1-\alpha)Q_i(s_t, p_{i,t}) + \alpha\left[\pi_{i,t} + \gamma \max_p Q_i(s_{t+1}, p)\right] \tag{1}$$

With this recursive method, agents can adjust their prices based on market feedback without human intervention. The hyperparameters used to produce such learning behavior have been tuned to replicate the actual high-frequency trading process. The value of the learning rate (α) was set to 0.125, and that of the discount factor (γ) to 0.95, thereby ensuring that the agents prioritize the stability of the market over short-term pricing. An ϵ-greedy policy has been used to regulate the exploration phase in all agent-based models considered here. The exploration parameter ϵ was initially set to 1.0 and then gradually decayed at a constant rate of $\beta = 1 - 10^{-5}$ per round, thereby allowing the market to evolve organically from a structure explorative phase to the stable and oscillatory phase depicted in the results.

The instantaneous reward or profit ($\pi_{i,t}$) for each agent is defined using the multinomial logit demand model. The following is the mathematical formulation of the aforementioned perfect substitutes assumption, along with the incorporation of demand uncertainty discussed earlier. The market share and demand of agent i can be calculated as shown in equation (2):

$$D_i = \frac{\exp((a - p_{i,t})/\mu)}{\sum_j \exp((a - p_{j,t})/\mu) + \exp(a_0/\mu)} \tag{2}$$

Here, a refers to the base level quality of the product, μ is used to represent the extent of horizontal differentiation (assigned a small value to ensure that the goods are close substitutes), and a_0 is the alternative available for those who do not buy. Finally, the stochastic demand noise was incorporated in the reward function by drawing Gaussian perturbations ($\epsilon \sim N(0, \sigma^2)$ with σ = 0.05), creating the precise environmental uncertainty that drives the volatility and convergence metrics analyzed throughout the study.

Figure 1 shows an emergent data structure created from the qualitative analysis of research conducted on autonomous AI agents v2.0. First-order concepts were organized into second-order concepts based on similarities via inductive analysis; these were then grouped into four longitudinal aggregate analytic dimensions.

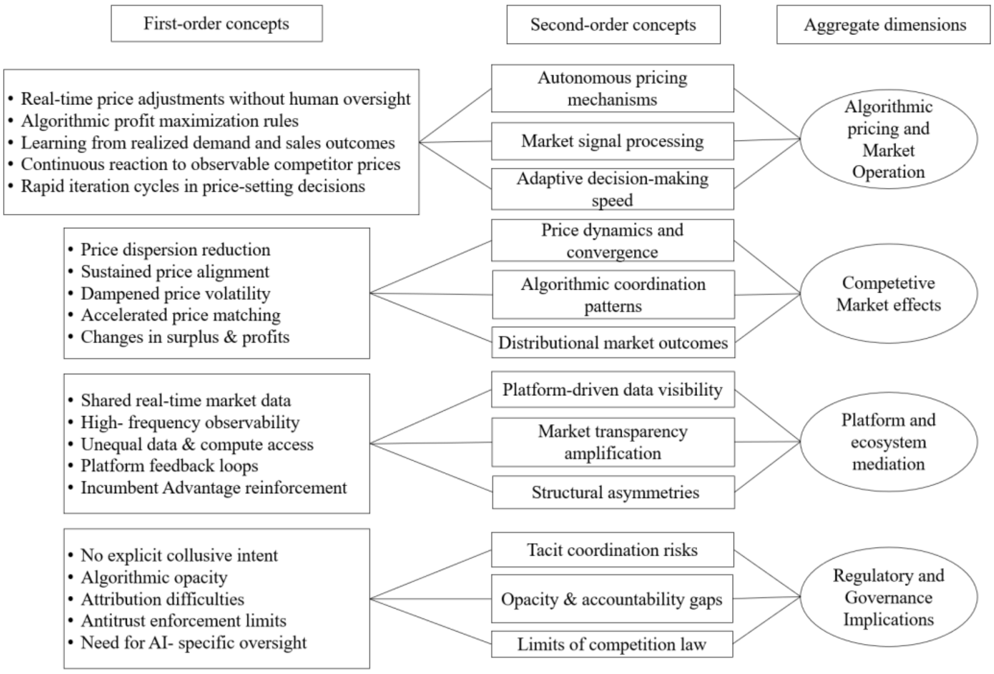

Fig. 1. Data structure of autonomous AI agents in digital markets

Aggregate Dimension 1: Algorithmic Pricing Capabilities and Market Operations. Autonomous AI agents' price without direct human involvement. In this aggregate dimension, elements of autonomous AIs pricing capabilities include the use of autonomous pricing, processing of market signals, and adapting decisions as a function of time. Autonomously priced products will reflect price changes continuously and instantaneously. This aggregate dimension illustrates various aspects related to real-time pricing changes, profit-maximizing algorithms, rapid iterative cycles, and adjustment to demand and sales activity based upon observable competitor pricing rather than through direct communication between autonomous AIs.

Aggregate Dimension 2: Competition's Effect on Market Prices. Market-level effects of autonomous AI agents on autonomous AI pricing activity. Aggregate dimension 2 incorporates both price movements and convergence, and the characteristics of the interaction among autonomous AIs. Examples of price movements/convergence include reduced price dispersion among sellers, the maintenance of price coordination among autonomous AIs over multiple interactions, decreased price volatility, faster overall price matching, and changes in consumer surplus and firms' profits.

The third aggregate dimension, termed Platform & Ecosystem Mediation, describes the mediation that digital platforms and market infrastructures have on the algorithmic behavior of market participants. The Platform & Ecosystem Mediation dimension captures the digital and algorithmic structures that are created within a market by way of their perceived benefits and uses, along with their ability to shape access to necessary information concerning digital and algorithmic market participation and behavior by market participants (i.e., such as the use of platform-mediated feedback loops, access to real-time data, and levels of price visibility).

The aggregate dimension entitled Regulatory & Governance Implications brings together the challenges that autonomous agents pose to pre-existing frameworks governing competition and regulation in the marketplace. The Regulatory and Governance dimension considers the potential for tacit collusion enabled by opaque algorithms and accountability gaps, and the limitations of existing competition laws as they relate to autonomous agents. The Regulatory and Governance Implications dimension therefore considers factors such as the absence of malicious intent on the part of multiple agents, the existence of tacit collusion through the algorithm used, the presence of algorithms that lack transparency, the difficulty of attribution for the algorithm, as well as the limitations of antitrust enforcement as applied to autonomous agents, and the need for the development and implementation of regulatory mechanisms specifically designed to address algorithmic dynamics in the marketplace.

An analytical tool for creating new government policy frameworks is provided, based on a data-driven approach that considers the degree of algorithmic independence and the availability of market information.

The proposal describes a range of autonomous algorithmic agents, demonstrating how each contributes to the emergence of new market forms, with varying degrees of predictive power and market adaptability.

The tool serves two main functions. It aims to create an analytical framework to examine how adaptive characteristics, such as learning speed and faster fulfillment of customer requests than competitors, will evolve and influence competitive outcomes. It also provides evidence to establish a formal link between observed competitive behavior and the corresponding required regulatory response. Essentially, the tool fills an analytical gap between policy frameworks at the micro-level (i.e., individual algorithms making decisions, how these decisions impact competition, and how regulators might respond) and at the macro-level (i.e., competition and regulation), as illustrated in Figure 2.

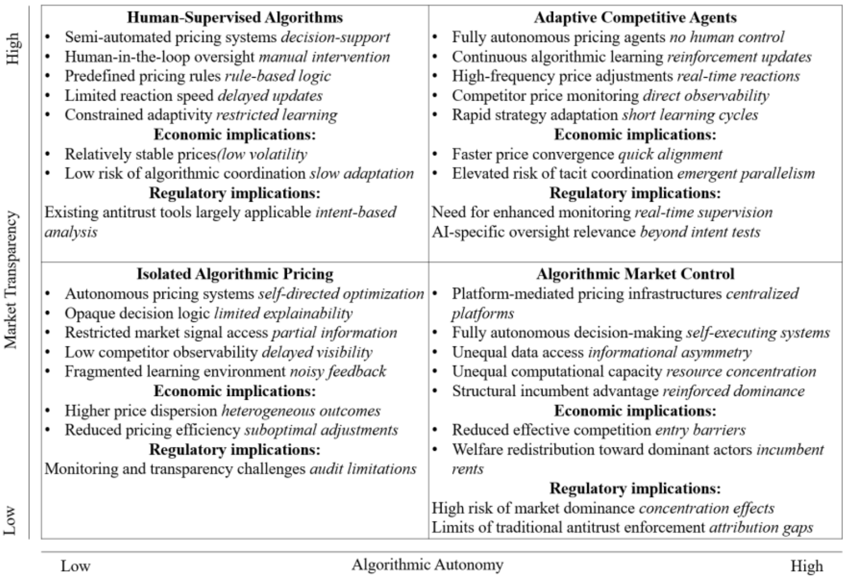

Fig. 2. Analytical framework of autonomous AI agents in digital markets

More specifically, as the study demonstrates, increasing algorithmic independence or the development of autonomous agents does not always lead to negative consequences for competition. For example, limited market transparency will lead to market fragmentation, as AI-based pricing agents will have little incentive to coordinate prices without direct market interdependence. Conversely, markets with very high levels of information transparency will further increase the interdependence of networks of autonomous algorithmic agents and are likely to lead to relatively rapid price alignment and implicit agreement even in situations where explicit collusion does not occur.

An equally important finding is the extent to which digital platforms fundamentally influence and reshape the fundamental structure of competitive markets. Simply put, digital platform pricing infrastructure can both increase the amount of available market information (through the underlying price comparison mechanisms hosted on these platforms) and simultaneously create significant structural inequalities in the market (due to agents' differential access to data and computing power). In turn, structural inequality leads to the development of market structures that exhibit significantly lower levels of effective competition and increased advantages for incumbents, demonstrating the ineffectiveness of existing intent-based antitrust policies.

The final aspect of this model will be explored through the analysis of empirical data collected through experimental comparisons. All experimental data will be used to classify empirical data into the four quadrants of the model to enable an assessment of the impact of autonomous algorithms on the efficiency and competitiveness of digital markets, as well as the effective incorporation of findings from experimental data into digital market governance.

Overall, the main goal of the model is to demonstrate that the duality of increased algorithmic independence and/or increased market transparency can lead to both increased market efficiency and the systemic vulnerabilities associated with this efficiency. By demonstrating the link between how AI algorithms make decisions and how these decisions will affect market structure and regulatory outcomes, this framework provides an analytical approach to interpreting empirical results and forms an analytically sound basis for assessing the extent to which existing competition policy and governance regimes will adequately address the issue of proper governance of AI in digital markets.

4. RESULTS AND DISCUSSION

This section examines the relationship between different autonomous pricing regimes and their effect on pricing dynamics, coordination patterns, profitability, and welfare outcomes. These differences among autonomous pricing models are being examined using a common simulation and comparison framework across human-supervised, adaptive-competitive, isolated, and platform-mediated control pricing regimes, in both their dynamic behavior and aggregate market outcomes.

Figure 3 is a summary comparison of the four analytical pricing regimes identified in the analytical framework. The analytical framework identifies distinct differences in their pricing behavior, efficiency, and market structure that arise from variations in algorithmic autonomy and transparency.

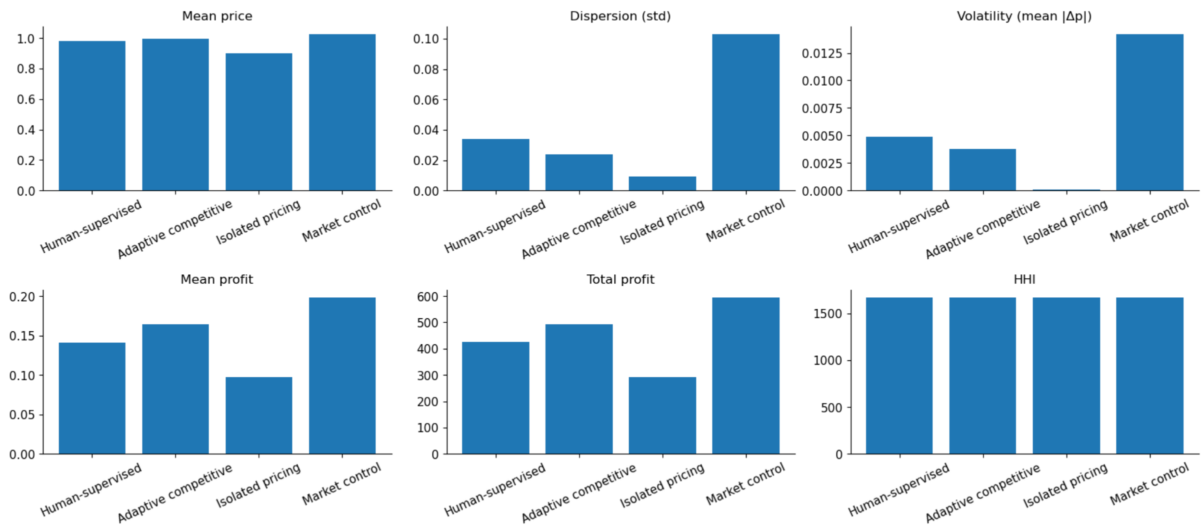

Fig. 3. Qualitative outcomes

The human-supervised pricing regime exhibits medium price levels, moderate dispersion, and relatively low volatility, which can be attributed to slower adjustment dynamics and continued human oversight of pricing. The adaptive competitive AI agents show higher average profitability and lower volatility in pricing behaviors, reflecting an increased ability to price efficiently and responsively. In addition, this increased ability to price efficiently and responsively has resulted in greater coordination or alignment of prices among sellers, as reflected in reduced dispersion of pricing. On the contrary, isolated algorithmic pricing (characterized by limited observability by competitors) led to lower average pricing and reduced profitability. However, it also exhibited fragmented and less efficient pricing. Lastly, algorithmic market control pricing yields the highest profit alongside high levels of price dispersion and volatility, thereby illustrating the emergence of structural asymmetries and concentration effects in highly autonomous/platform-mediated environments. Ultimately, the observed relationship demonstrates that increased algorithmic autonomy does not necessarily create competitive harm, as it depends on how autonomy interacts with transparency in generating market outcomes, which supports the proposed framework presented in Figure 2. Moreover, although adaptive AI pricing may be better at producing efficiencies than others, there will be situations where they increase the risk of tacitly coordinating or dominating a market, leading to concerns that go beyond conventional antitrust assessments of intent.

To provide a comprehensive comparison of autonomous pricing regimes, the simulation results are summarized using a multi-dimensional analytical framework that evaluates price dynamics, dispersion, profitability, and welfare outcomes across different levels of algorithmic autonomy and market transparency. The comparative outcomes highlight structural differences between human-supervised, adaptive autonomous, isolated, and platform-mediated pricing systems. These differences reflect how algorithmic independence and information availability influence convergence behavior, market stability, and economic efficiency. The integrated comparison of these regimes is presented in Figure 4.

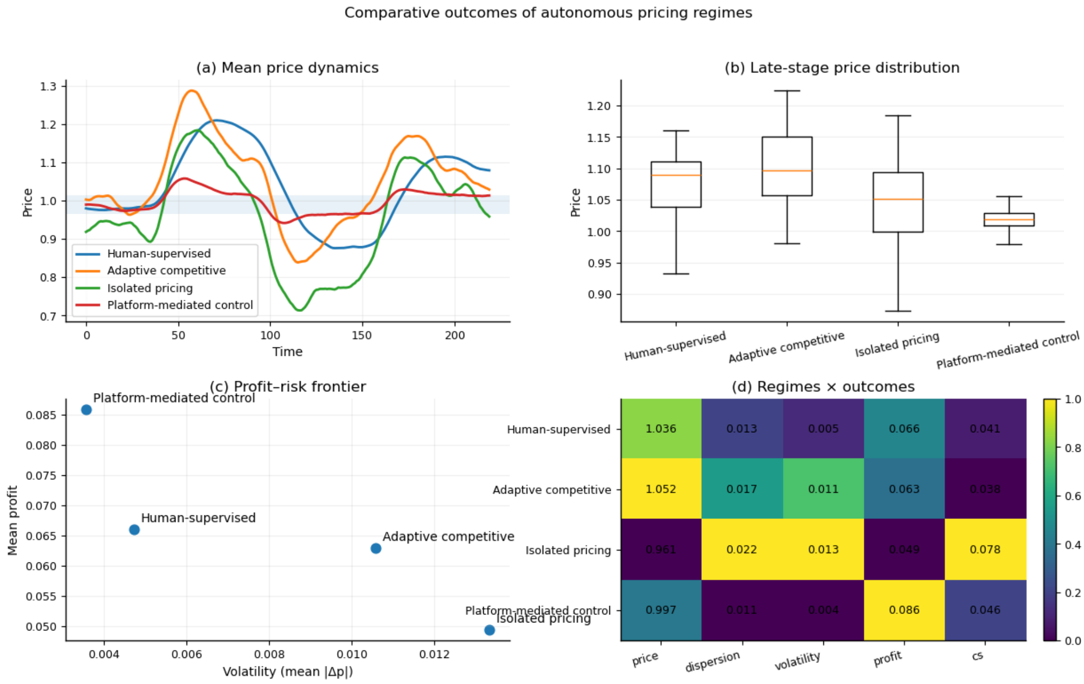

Fig. 4. Comparative outcomes of autonomous pricing regimes

The smoothed time series of prices for three categories of autonomous pricing regimes is shown in Figure 4(a). It may be noted that the price adjustments made by the human-supervised pricing regime are more dampened than those made by the adaptive competitive regime, as they are less responsive to exogenous shocks, while the price adjustments made by the adaptive competitive agents are more rapid and occasionally overshoot their target prices. The isolated pricing regime exhibits higher stochastic noise and lower convergence rates than the other two regimes, while the platform-mediated control regime maintains prices within a constrained range and allows limited experimentation in an adaptive environment.

Figure 4(b) presents the price distributions of individual sellers in the late stage, showing significant price dispersion across the two autonomous pricing regimes of isolated and adaptive competitors. On the contrary, the prices of sellers supervised by a human exhibit tight price clustering, whereas those of sellers under platform-mediated control regimes are clustered along a continuum, with controlled dispersion consistent with pricing discrimination rather than regulatory stabilization.

Figure 4(c) presents the profit-risk trade-off at the extremes of the price-risk trade-off; in particular, there exists a trade-off between average profitability and price volatility. While adaptive competitors achieve higher average profits than human-supervised pricing, they do so at the cost of greater price volatility. Conversely, the pricing regime of human supervision is more oriented towards stability than to profit maximization. The pricing regime of platform-mediated control lies between the other two pricing regimes, yielding profitability averages greater than those of either regime while exhibiting moderated risk exposure.

Figure 4(d) is the normalized outcome matrix for the analyzed pricing regimes. In looking at price level, dispersion, volatility, profit, and consumer surplus, all three pricing regimes had distinct characteristics across these metrics. Therefore, it can be stated that no single pricing regime dominates the other two across all metrics. Instead, the results highlight the relative structural trade-offs between efficiency, stability, and welfare across autonomous digital markets.

The differences in price convergence and overall dispersion at the end of pricing sessions across four price-setting systems are presented in Figure 5. The methods we used to measure price convergence between price systems are convergence envelopes constructed from the following two indicators: (1) a price trajectory based on the median price, and (2) a price trajectory based on the interquartile range (IQR).

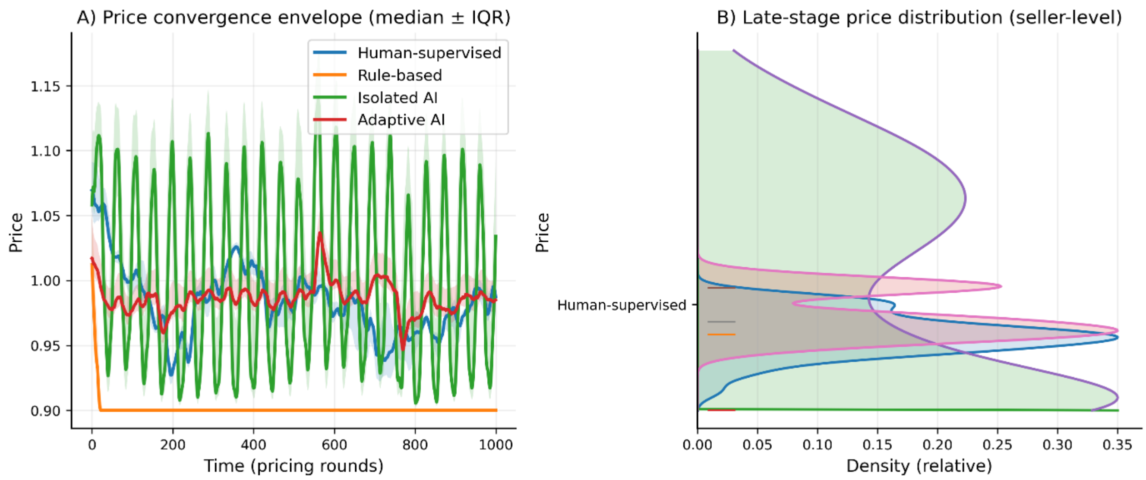

Fig. 5. Price convergence and end-state distributions

As shown in Panel A, there is strong convergence in human-supervised pricing toward a single price level at or near the stability zone (0.98 and above) within approximately 150-200 pricing rounds, with a very narrow interquartile range of about 0.02. Therefore, an effective means of suppressing short-term price fluctuations via supervision was utilized. Furthermore, the rules-based price setting experiences a quick decline and converges to the lowest price range (0.90). In comparison, isolated AI pricing shows persistent oscillations throughout the simulation period. The median price fluctuated over a much larger price band (0.93-1.08), and the associated IQR is also quite large (greater than 0.06), suggesting that these differentially dispersed prices never achieved convergence (even after 1000 rounds of pricing) due to the fragmented learning and lack of coordination between autonomous agents.

Adaptive AI pricing initially converges faster to stable prices than isolated AI pricing; however, it settles into the equilibrium range after about 100 rounds of pricing. That said, adaptive AI pricing has an IQR about 50% larger than that of human-supervised pricing (IQR for adaptive AI pricing: 0.03-0.04), and several intermittent overshooting tails exist, particularly around midpoint price fluctuations. Therefore, while adaptive autonomy promotes faster adjustment to equilibrium prices, it cannot eliminate intrinsically unstable prices.

Panel B provides additional confirmation of the dynamic findings through late-stage price distributions whose results were based on the price received at each seller level. Human-supervised pricing yields compactly distributed, unimodal price distributions, with prices concentrated around the equilibrium price and relatively little price tail mass. Thus, adaptive AI pricing exhibits a mildly asymmetric price distribution (broader than the human-supervised price distribution), suggesting that sellers continue to experiment with prices. In contrast, isolated AI pricing generated the most dispersed prices (i.e., prices varied more broadly than the other two pricing methods and exhibited persistence in this variability even during stationary time periods that represented nearly the entire price range (0.90-1.15)). The last result was that the rule-based baseline was strictly confined to the lower price threshold. In summary, the speed of price convergence and the quality of that convergence are separate dimensions of market stability. Adaptive autonomy, while responsive to changing conditions, does not enable sustained price coordination or the reduction of dispersed prices across pricing systems without incorporating some form of supervisory or coordinating mechanism.

Figure 6 presents an analysis of the robustness and sensitivity of autonomous pricing regimes against key market characteristics, including demand elasticity, stochastic demand noise, and scalability with the number of agents. The results in Panel (a) demonstrate an extremely high sensitivity of profitability to changes in elasticity between b=1.2 and b=2.4 across the different pricing regimes. As average elasticity increases, average profit falls dramatically, approaching zero across all pricing methods. However, for average elasticities in the moderate range, approximate b=1.8, profitability levels of the adaptive competitive and platform-mediated control regimes are significantly higher than those of the isolated pricing regime (0.08-0.09 and 0.05), demonstrating these pricing methods have a greater degree of resiliency to competitive forces.

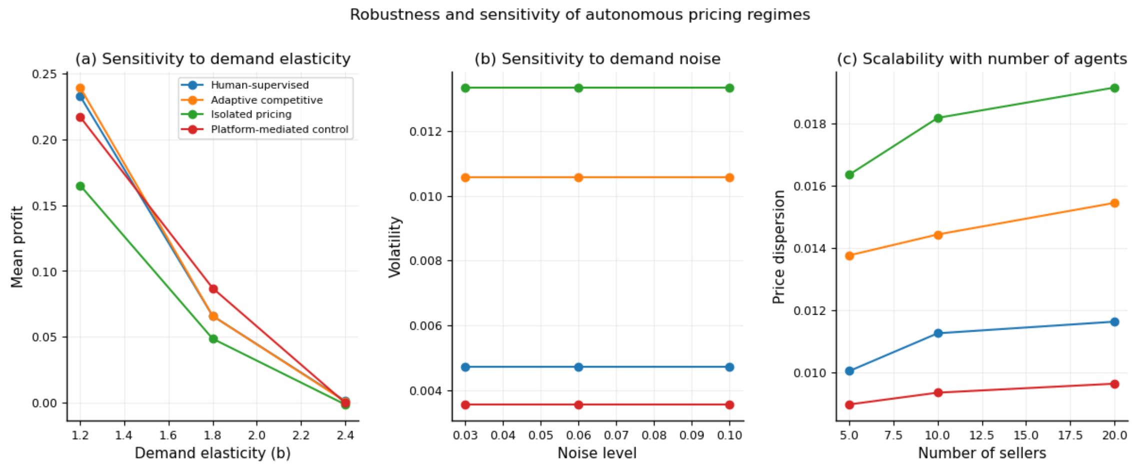

Fig. 6. Robustness and sensitivity of autonomous pricing regimes

The objective of Panel (b) was to assess the sensitivity of the different autonomous pricing regimes to stochastic demand noise. Overall, there was little variance in volatility across the different regimes within the range of stochastic demand noise 0.03-0.10; however, aggregate volatility across the different regimes was significant. Average levels of volatility associated with isolated pricing (0.013) were disproportionate to those associated with human-supervised pricing (0.005), by a factor of almost three. Adaptive competitive pricing was intermediate in volatility (0.010), while platform-mediated control pricing had the lowest volatility, 0.003-0.004. Thus, indicating that platform-mediated control pricing is the most robust to stochastic perturbations.

Panel (c) assessed the impact of scale by analyzing price dispersion within each pricing method/regime in terms of the number of competing sellers, ranging from 5 to 20 sellers. All pricing methods/regimes showed increased price dispersion as the number of competing sellers increased; however, the rate at which dispersion increased varied significantly across methods/regimes. Isolated pricing shows the steepest growth in price dispersion, from approximately 0.016 to nearly 0.019, whereas adaptive competitive pricing exhibits a much slower growth in dispersion, from 0.014 to 0.0155. Human-supervised pricing had the lowest increase in dispersion, rising from 0.010 to 0.0115; furthermore, Platform-mediated control exhibits the lowest dispersion across all levels of complexity (less than 0.010), thus demonstrating a higher level of scalability when faced with increased market complexity.

The overall robustness analysis demonstrates that qualitative rankings of the pricing regimes will be stable across a broad spectrum of market conditions. The differences in price stability, average profitability, and average levels of coordination observed among the pricing regimes are not simply functions of the specific absolute values assigned to individual parameters but rather reflect inherent structural attributes associated with each pricing method/regime concerning autonomy, transparency, and information asymmetry.

While our model provides clear insight into the behavior of pricing algorithms, there are important limitations to consider when extrapolating our results to realistic situations. For instance, our assumption that all products are perfect substitutes implies a much higher rate of coordination than would be achieved in practice. Indeed, due to the effects of loyalty and product differentiation, the actual speed of price coordination and the likelihood of collusion might be much slower than our findings suggest. Moreover, we idealized the model and set the marginal costs equal across all agents, simplifying the complexity of a real-world scenario with constantly changing economies of scale, constrained warehouse capacity, and unpredictable supply shocks. The existence of different variable cost levels among our retailers will lead to deviations from the observed stable price dynamics. Lastly, in contrast to our model, price updates in a real digital economy occur at different points in time. Network lags inevitably introduce noise into the system, hindering algorithmic price coordination.

5. CONCLUSIONS

Digital marketplaces increasingly rely on autonomous artificial intelligence agents to determine prices through their behavior and to affect the level of competition and welfare through their interactions with each other, rather than on firms' decisions. As a result, the dynamics of these marketplaces are shaped by a continuously evolving ecosystem based on three key characteristics: algorithmic autonomy, information transparency, and platform mediation. Together, these characteristics will shape the nature of competition in AI-generated environments.

Each of these three essential elements of the marketplace can be configured to yield numerous types of economic trade-offs. For example, markets that are heavily reliant on human oversight and have strict price limits tend to produce stable, low-coordination-risk environments; conversely, a market that is highly autonomous and/or heavily mediated by platforms will produce very high levels of efficiency/profitability while increasing both concentration and governance-related challenges. As such, algorithmic design alternatives and market governance considerations must be viewed as interdependent. The framework used to address competition and regulatory issues in digital marketplaces needs to reflect the interconnectedness of these two elements rather than rely on a traditional intent-based approach.

For implementation, regulation should stop requiring proof of explicit collusion and instead enforce structural “compliance by design” mechanisms. First, competition regulators should impose mandatory pre-market algorithmic sandboxing on platforms to prove that their pricing agents do not enter into implicit collusive equilibria in simulated stress tests before entering the market. To address algorithmic black boxes, this should be accompanied by dynamic audit trails. Competition regulators should require digital platforms to log their algorithms’ reward functions, learning parameters, and live data feeds in industry-standard, cryptographically secure formats to enable automated post-hoc antitrust audits. Finally, to mitigate the structural power of the platform's algorithmic control, policymakers should enforce API-driven information parity. This will force digital platforms to provide third-party sellers equal access to the high-frequency demand signals and price histories used by incumbents’ algorithms.

Future research could expand upon the current framework by introducing product heterogeneity, asymmetric cost structures, and more highly nonlinear pricing functions. Other directions for future research could focus on alternative learning architectures (such as rule-based versus hybrid reinforcement-based learners) and the influence of various information-sharing delays and/or information asymmetries on decision-making processes. Additionally, the continued expansion of this framework to larger market sizes and more complex platform governance systems will produce greater clarity regarding scalability and concentration impacts. Finally, empirical validation of this framework using actual pricing datasets would further confirm the link between simulation-based insights and practical marketplace outcomes, thereby providing a basis for developing more accurate regulatory assessments of the use of autonomous AI agents in digital marketplaces.

Conflicts of Interest

The authors report there are no competing interests to declare.

REFERENCES

Agrawal, A., Gans, J. S., & Goldfarb, A. (2019). Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, *33*(2), 31–50. https://doi.org/10.1257/jep.33.2.31

Bajari, P., Chernozhukov, V., Hortaçsu, A., & Suzuki, J. (2019). The impact of big data on firm performance: An empirical investigation. *AEA Papers and Proceedings*, *109*, 33–37. https://doi.org/10.1257/pandp.20191000

Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, *110*(10), 3267–3297. https://doi.org/10.1257/aer.20190623

Comunale, M., & Manera, A. (2024). The economic impacts and regulation of AI: A review of the academic literature and policy actions. *IMF Working Papers*, *2024*(065), 1. https://doi.org/10.5089/9798400268588.001

Gerpott, T. J., & Berends, J. (2022). Competitive pricing on online markets: A literature review. *Journal of Revenue and Pricing Management*, *21*(6), 596–622. https://doi.org/10.1057/s41272-022-00390-x

Goldfarb, A., & Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, *57*(1), 3–43. https://doi.org/10.1257/jel.20171452

Groeneveld, J., Herrmann, J., Mollenhauer, N., Dreeßen, L., Bessin, N., Tast, J. S., Kastius, A., Huegle, J., & Schlosser, R. (2024). Self-learning agents for recommerce markets. *Business & Information Systems Engineering*, *66*(4), 441–463. https://doi.org/10.1007/s12599-023-00841-8

Harrington, J. E. (2018). Developing competition law for collusion by autonomous artificial agents. *Journal of Competition Law & Economics*, *14*(3), 331–363. https://doi.org/10.1093/joclec/nhy016

Kastius, A., & Schlosser, R. (2022). Dynamic pricing under competition using reinforcement learning. *Journal of Revenue and Pricing Management*, *21*(1), 50–63. https://doi.org/10.1057/s41272-021-00285-3

Klein, T. (2021). Autonomous algorithmic collusion: Q-learning under sequential pricing. *The RAND Journal of Economics*, *52*(3), 538–558. https://doi.org/10.1111/1756-2171.12383

Krarup, T., & Horst, M. (2023). European artificial intelligence policy as digital single market making. *Big Data & Society*, *10*(1), 20539517231153811. https://doi.org/10.1177/20539517231153811

Lancieri, F., Edelson, L., & Bechtold, S. (2025). AI regulation: Competition, arbitrage and regulatory capture. *Theoretical Inquiries in Law*, *26*(1), 239–262. https://doi.org/10.1515/til-2025-0010

Li, C., & Virtosu, I. (2025). Competition law in the age of AI: Confronting algorithmic collusion in the smart economy. *Smart Cities and Regional Development (SCRD) Journal*, *9*(2), 41–54. https://doi.org/10.25019/qnjgmk44

Li, Q. (2025). The economics of AI-enabled price discrimination. In *AI-enabled Price Discrimination* (pp. 1–37). Springer. https://doi.org/10.1007/978-3-031-84790-5_1

Normann, H.-T., & Sternberg, M. (2023). Human-algorithm interaction: Algorithmic pricing in hybrid laboratory markets. *European Economic Review*, *152*, 104347. https://doi.org/10.1016/j.euroecorev.2022.104347

Schauer, A., & Schnurr, D. (2023). Competition and collaboration between human and artificial intelligence in digital markets. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4565678

Shahzad, M. H., Rafy, T., Malik, A., Waqar, M., & Ahmad, S. (2025). The rise of autonomous enterprises: How AI is reshaping business models and market dynamics. *ACADEMIA International Journal for Social Sciences*, *4*(2), 677–693. https://doi.org/10.63056/ACAD.004.02.0207

Spulbar, L.-F. (2025). Legal frameworks for AI-driven markets and their challenges and opportunities in the digital economy. *Revista de Ştiinţe Politice. Revue des Sciences Politiques*, (86), 288–304.

Wu, J. J.-X. (2023). Beyond free markets and consumer autonomy: Rethinking consumer financial protection in the age of artificial intelligence. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4566590

Yavuz, T., & Kaya, O. (2024). Deep reinforcement learning algorithms for dynamic pricing and inventory management of perishable products. *Applied Soft Computing*, *163*, 111864. https://doi.org/10.1016/j.asoc.2024.111864