OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

Algorithmic Pricing and AI-Driven Competitive Strategies: A Game Theory Review

Vedang Tiwari
vedang2125@gmail.com

ABSTRACT

Artificial Intelligence (AI) has revolutionized digital markets and pricing by empowering companies to leverage machine learning, reinforcement learning, and predictive analysis to automate strategic pricing decisions. Algorithmic pricing enhances efficiency, personalization and responsiveness but also poses risks of tacit collusion, anti-competitive coordination and regulatory shortcomings. In this review, the literature published from 2015 to 2025, related to algorithmic pricing and its game-theoretic perspective, is critically synthesised as a narrative review. The paper explores the interactions between AI systems in repeated strategic settings, the convergence of reinforcement learning agents towards collusive equilibria, and the amplification of algorithmic coordination in platform-based markets. This paper brings together game theory, AI governance, industrial organization and digital platform economics, as opposed to prior reviews such as Ezrachi and Stucke (2017) and the OECD (2017), which are more focused on either economic efficiency or antitrust implications. The review reveals that, despite empirical work, the literature does not have a satisfactory theoretical and policy understanding of how an algorithm can autonomously reach collusive equilibria, nor how it can be regulated by laws that assume intentionality in humans, thus making them inappropriate to govern the coordination that is generated by algorithms. To fill this void, the paper suggests the notion of ‘emergent strategic coordination’—that is, market coordination that is computationally produced through self-learning optimization processes in the absence of intentional collusion, a concept that builds on, rather than supplants, existing frameworks developed by Harrington (2018) and Ezrachi and Stucke (2017). The paper suggests that current antitrust laws have failed to keep up with AI’s strategic interactions, especially since they were drafted for humans rather than AI systems. The paper concludes that algorithmic pricing is much more than a new technology in business practices, it is a transformation in the nature of competition and the requirement for new business theory and regulatory frameworks.

Keywords: Algorithmic Pricing; Artificial Intelligence; Competition; Game theory; Strategic Pricing

1. INTRODUCTION

Markets are now fast being digitalized, and this has changed the nature of competition in contemporary business settings. Artificial intelligence (AI) and algorithmic pricing systems are becoming the new norm for firms looking to fine-tune prices as they go to meet the demand of consumers, the actions of competitors, and the whims of the market. Autonomous pricing algorithms are widely used in various industries, including e-commerce, ride-sharing, travel, food delivery, online advertising, and financial trading, to optimize pricing strategies and improve market efficiency (Calvano et al., 2020). Dynamic pricing systems powered by AI-based strategic pricing are most visibly deployed by companies such as Amazon, Uber, and Airbnb.

Algorithmic pricing has garnered a considerable amount of academic research as pricing algorithms can learn and adapt continuously. AI systems can analyze vast amounts of market information in real-time and act quickly to market changes, unlike the traditional human-centric pricing models. The ability to carry out these practices can impact the level of competition in the market, as they can allow firms to implicitly coordinate prices without explicit communication (Ezrachi & Stucke, 2017). These things have raised even greater awareness of potential collusion, anti-competitive practices and the effectiveness of the current antitrust system.

An important theoretical approach to the understanding of these developments is the Game theory approach as pricing decisions in this market inherently contain strategic interdependence. In classical Bertrand competition, price competition will drive firms down to the marginal cost. Firms can achieve supra-competitive prices through collusion by consistently playing the same game through repeated-game theory, but with algorithms increasingly controlling pricing, such collusion becomes much easier to execute as the algorithms respond to the prices charged by competitors and adjust their prices based upon those responses (Tirole, 2015).

Recent evidence suggests that reinforcement-learning algorithms can also learn to collude without any overt communication between the firms in repeated pricing games (Calvano et al., 2020; Klein, 2021). This raises significant concerns regarding antitrust enforcement because traditional antitrust law is predicated on the notion that proof of intent, communication, or agreement is required to address collusion. Since algorithmic coordination is independent and autonomous, it is much more difficult to interpret collusive behaviour under existing antitrust law as algorithmic pricing continues to change competitive dynamics across industries.

The rise of AI in competitive pricing coincides with other broader structural changes affecting platform economies. Digital platforms generate massive volumes of behavioural data; create network effects; and centralise the market's visibility as a result of algorithmic agents operation, thus enhancing the degree of strategic interaction among all of the various algorithmic agents (Parker et al., 2016). Therefore, algorithmic pricing is no longer simply a technological advancement; rather, it represents an institutional evolution that dramatically alters the competitive dynamics associated with many different types of industries.

The literature reviewed critically addresses the issue of algorithmic pricing and competitive strategy based on Artificial Intelligence (AI) from a game-theoretic perspective, taking into account the period between 2015 and 2025. Drawing on economics, AI governance, industrial organization and competition law, the paper examines the impact of AI systems on strategic interaction in digital markets. One important point that has been observed, and is the basis for this review, is that empirical scholarship has advanced significantly faster than regulatory and theoretical scholarship in regulating markets that reach outcomes different from competitive ones without the direct involvement of human agents. Algorithms have been demonstrated to give rise to collusive outcomes, as in Calvano et al. (2020) and Klein (2021), without any direct human intervention, yet the regulatory and theoretical frameworks intended to govern them have failed to keep up. While the literature on algorithmic pricing focuses on the regulatory or technical aspects, this paper provides a conceptual framework for understanding emergent strategic coordination between algorithms under self-learning, and contends that novel theoretical and policy tools are needed to understand and address it.

2. METHODOLOGY

This paper takes the form of a narrative literature review rather than a systematic review. Sources were identified through targeted searches of Google Scholar, SSRN, and the websites of major competition authorities (including the OECD and the UK Competition and Markets Authority), using search terms including ‘algorithmic pricing’, ‘algorithmic collusion’, ‘AI and antitrust’, ‘reinforcement learning collusion’, and ‘platform pricing competition’. TIt is based on around 30 published sources from 2015 to 2025, including in the field of industrial organisation, AI and machine learning, digital platform economics and competition law. The starting point for 2015 is where empirical and computational research on algorithmic collusion started to emerge in the economics literature (with the exception of the works by Fudenberg and Tirole (1991) on repeated game theory and Sutton and Barto (2018) on reinforcement learning, which were cited as background theoretical literature, but not included in the reviewed literature on algorithmic pricing itself).

The paper is a narrative review, and does not use a prespecified method of searching databases and does not have a set of inclusion and exclusion criteria or multiple reviewers to independently screen the articles. Instead, source relevance to the core themes of the paper (the game-theoretic foundations of algorithmic pricing, the empirical and computational evidence of collusion, platform-level structural factors, and regulatory responses) and citation prominence and publication venue were used to select sources. An aim that is conceptual and synthetic, as in the present paper, is an appropriate approach. It also has a limitation, as stated in Section 8, in that the corpus is not a result of a search procedure that can be reproduced; the coverage statements in the paper should be interpreted with this limitation.

This literature is evolving rapidly: three of the three sources referred to (Assad et al., 2020; Fish et al., 2024; Miklos-Thal & Tucker, 2024) are working papers or preprints and are not peer-reviewed at the time of writing. The paper stands in opposition to two streams of review literature: studies that primarily focus on the economic efficiency controversy surrounding algorithmic pricing (e.g. Brown & MacKay, 2023); and studies that primarily focus on antitrust and regulatory issues (e.g. Ezrachi & Stucke, 2017; OECD, 2017). This paper also argues that these should be read together as they rely on the same technology one way or the other: algorithmic pricing offers efficiency gains, while it offers coordination risks as well. This is elaborated more in Section 7.

3. GAME THEORY AND COMPETITIVE PRICING

It is a well-established practice to model competitive behaviour in oligopoly markets with the game theory. Pricing competition is a situation in which companies strategically determine their prices in light of the likely reactions of their competitors. The Bertrand model assumes firms compete by setting prices simultaneously, leading to rational firms to set prices at the level of the marginal cost (Belleflamme & Peitz, 2021). Repeated-game theory adds to this picture, as repeated interaction can give firms opportunities to coordinate tacitly (Fudenberg & Tirole, 1991): If repeated interaction and the possibility of continued high profits from cooperation in the short term are included, then cooperation over time can be supported. This is a form of collusion between firms that occurs when each firm sees an interest in maintaining the prices above competitive levels, but does not actually communicate directly. A balance of this kind has been achieved in the past by maintaining equilibrium in the markets, transparency of prices, and known behavior from competitors.

Algorithmic pricing systems are expected to improve the transparency of the markets and to lessen the strategic uncertainty of the tacit collusion described above. These enable companies to watch the competition's prices at all times, to observe consumer habits and work out demand and market trends, and to respond quickly to price changes. Brown and MacKay (2023) conclude that algorithmic pricing introduces a new competitive landscape, enabling firms to respond to it strategically more rapidly and successfully than previously. Reinforcement learning is especially relevant here: reinforcement learning agents make their decisions based on their experience with their environment: by repeatedly interacting with the environment to maximise cumulative reward over time (Sutton & Barto, 2018), and in a pricing context, this involves observing the actions of their competitors, calculating the profits as a result of the observed actions, and adjusting the actions accordingly, similar to the logic of repeated games in classical theory.

We showed in a computational simulation that independent Q-learning algorithms have a tendency to arrive at collusive pricing equilibria (Calvano et al., 2020). In fact, their simulations revealed that reinforcement learning agents could learn to charge supra-competitive prices, as well as punish agents who deviated from this in their behavior. The discovery was of importance to economists and regulators because it showed that algorithmic collusion can be a side-effect of the optimisation process – no one intended it, or even knew it was happening. This is an important point to reiterate, since it is not a real market result, but a simulation outcome; this is discussed further below and in Section 8, and is an important distinction regarding the influence of this result in the research and policy discussion.

It is important that implications of this finding are significant exactly because algorithmic coordination of this kind is different from collusion, as it is traditionally understood. Usually, collusion between human actors requires some type of deliberate communication or meeting. Algorithmic coordination can, on the other hand, occur without the firms using the algorithms engaging in any deliberate collusion in the adaptive learning process. Klein (2021) similarly found, using a Q-learning model under sequential pricing, that self-learning algorithms can sustain stable cooperative pricing equilibria; this is a finding about the behaviour of the algorithmic system studied, and Klein's paper does not compare algorithmic and human decision-making directly.

These findings point to a more general question about whether the assumptions underpinning classical competition theory remain adequate once pricing decisions are delegated to algorithms. Traditional theories of competition assume bounded rationality and imperfect monitoring on the part of market participants. Algorithmic systems, by contrast, can process information more completely and respond more quickly and consistently than human decision-makers. This raises the possibility that some classical assumptions about market equilibrium may need revisiting in digital pricing environments, although, as discussed in Section 8, the extent to which simulation-based findings of this kind generalise to real, more complex markets remains an open empirical question.

4. PRICING IN PLATFORM ECONOMIES: THE ROLE OF AI

Digital platforms are a particularly significant setting for the growth of algorithmic pricing. Platform firms like Amazon, Uber, Airbnb and Booking.com employ dynamic pricing models that can continuously and automatically alter prices as conditions in the market evolve. These systems rely on machine learning, predictive analytics and consumer behaviour data to make real-time dynamic pricing adjustments.

Platform economies's sharpen the strategic interaction by providing market information and greater price transparency. Consumers can easily make comparisons between the prices of different sellers almost instantaneously, while businesses can closely track competitors' activities at all times. Parker et al. (2016) claim that platform ecosystems create a meaningful network effect that can affect competition, as well as an increasing tendency for firms in these ecosystems to focus on strategic positioning instead of price as a way to compete.

Two areas where dynamic pricing has gained a particular importance are ride-sharing and ecommerce. The surge pricing mechanism dynamically adjusts pricing depending on the imbalance between supply and demand, and, according to Hall et al. (2015), such a system can help to keep drivers available and minimize passenger waiting times; they report on one situation where the system was briefly disabled, resulting in a dramatic rise in waiting times for passengers, highlighting the normal operation of the system. Similarly, Amazon has automated pricing algorithms that continuously adjust to the prices of its competitors and consumer demand (Chen et al., 2016). The efficiency and responsiveness of these systems is desirable, but also provides more potential for algorithmic coordination.

Deng (2018) points to a number of characteristics of algorithmic agents that could make them more likely than humans to be collusion-friendly: Algorithms make faster decisions based on their opponents' actions, remember more of what has gone on previously, and continuously watch for deviations. For this reason, in markets that are extremely fast moving, algorithms may be able to maintain a supra-competitive equilibrium more effectively than the human manager.

Another key economic aspect recognized in the platform economics literature is data asymmetry. Having vast amounts of consumer data allows companies a competitive edge—predictive analytics relying on more comprehensive data sets can help optimize pricing and forecast demand (Varian, 2019). This may, therefore, lead to market power and anti-competitive dynamics generally in the hands of a few big platforms.

The rise of new AI-driven platforms has also changed the structural challenges for new players. More powerful computing systems, and better data, in the hands of dominant firms allows them to use more complex pricing algorithms, adding to their competitive edge. This has led to concerns about digital monopolisation and the exclusionary behaviour (Autorité de la concurrence & Bundeskartellamt, 2019).

In fact, digital platforms have become a strategic ecosystem where pricing is no longer an isolated function but part of the whole platform and is entangled in the overall platform strategy, with recommendation systems and personalisation technologies becoming more and more crucial to digital pricing. As far as the present author could tell from the literature, the importance of these structural interactions has received a bit less attention than has been devoted to the pricing algorithms themselves.

5. REINFORCEMENT LEARNING AND ALGORITHMIC COLLUSION

Table 1: Summary of key studies included in the review

Study	Method	Setting	Key Findings
Calvano et al. (2020)	Computational simulation (Q-learning)	Simulated oligopoly	Independent reinforcement learning agents converge on supra-competitive prices and punish deviations, with no explicit communication.
Klein (2021)	Computational simulation (Q-learning, sequential pricing)	Simulated oligopoly	Self-learning algorithms sustain stable cooperative pricing equilibria.

Lepore (2021)	Computational simulation (Q-learning)	Simulated market with volatility and incomplete information	Collusive outcomes persist even under volatility and incomplete information, suggesting the result is not an artefact of unusually stable conditions.
Assad et al. (2020)	Empirical (real-market)	German retail petrol stations	Adoption of algorithmic pricing software is associated with reduced competition, providing a real-world echo of the simulation findings.
Werden & Froeb (2018)	Conceptual/critical commentary	General claims of market concentration	Cautions that striking simulation or stylised results do not necessarily indicate a first-order problem in real, more complex markets.
Miklos-Thal & Tucker (2024)	Theoretical/computational	Reinforcement learning pricing models	Under some conditions, RL systems produce unstable or chaotic dynamics rather than stable collusion, broadening the risk profile of algorithmic pricing.
Fish et al. (2024)	Computational simulation (LLM agents)	Strategic pricing interactions	LLM-based agents can engage in collusive pricing autonomously, extending coordination risk beyond reinforcement-learning architectures.
Hall et al. (2015)	Empirical (real-market)	Uber surge pricing	Surge pricing maintains driver availability and reduces passenger waiting times; a system outage caused a sharp rise in wait times.
Brown & MacKay (2023)	Empirical/theoretical	Algorithmic pricing markets	Price transparency can constrain as well as enable coordination, since it allows competitors and new entrants to undercut collusive prices.

The term reinforcement learning is a common term in the literature on algorithmic collusion. Reinforcement learning systems are different to rule-based algorithms because they learn from their interaction with an environment, and seek to maximize the total reward over a long period of time instead of solely executing pre-programmed instructions. This is an effect that was demonstrated by Calvano et al., 2020 (see above), who found that independent reinforcement learning agents repeatedly arrived at supra-competitive pricing equilibria in simulated oligopolistic markets without any explicit communication taking place.

It is however a much more open question whether this simulation result is what would be observed in real markets and the literature does not agree on this. Lepore (2021) builds on this simple result, demonstrating that even under conditions of volatility and incomplete information, reinforcement learning agents can achieve collusive results, thus showing that it is not just a byproduct of a very stable, fully observed simulation environment. By contrast, the results of Assad et al. (2020), a real market study of the adoption of algorithmic pricing software in the market for petrol stations in Germany, suggests that the patterns seen from the simulation results do extend to the real market, with a decrease in competition after adoption. Werden and Froeb (2018) make a more skeptical point, that it would be premature to assume that market concentration is a problem of first order in the current market, and, by extension, that algorithmic coordination is a problem of first order in the current market. These three sources seen together indicate that this basic lab finding is reasonably stable across variations in experimental conditions and at least has a real-world echo, but that the practical significance of algorithmic collusion in the outside world isn't a foregone conclusion.

A further difficulty with classical game-theoretic models is the possibility for instabilities to arise in competitive markets through reinforcement learning. Miklos-Thal and Tucker (2024) show that, under some conditions, reinforcement learning systems can produce unstable or chaotic market dynamics rather than settling into a stable collusive equilibrium, suggesting that algorithmic pricing may increase volatility as well as the risk of supra-competitive pricing.

A further obstacle to regulating reinforcement learning systems is their opacity. Deep learning models of the kind used in many pricing algorithms are frequently described as ‘black box’ systems, in the sense that it is difficult for regulators, and sometimes even for the deploying firms themselves, to establish why a particular pricing strategy emerged (Burrell, 2016). This opacity creates considerable legal uncertainty regarding liability for collusive outcomes produced by such systems.

A more recent strand of research has begun to examine the use of large language models (LLMs) in pricing strategy. Fish et al. (2024) show, in a series of strategic interactions, that LLM-based agents can engage in collusive pricing behaviour autonomously. The contextual reasoning and linguistic flexibility of LLMs may broaden the range of mechanisms through which algorithmic coordination can arise, beyond the reinforcement-learning architectures discussed above; in this setting, strategic behaviour is shaped by factors such as prompt design and contextual framing, which existing competition frameworks are not well equipped to address.

6. REGULATORY ISSUES AND ANTITRUST LIMITATIONS

The current antitrust regimes have serious shortcomings when it comes to algorithmic pricing. Past theories of competition law have been very much dependent on the existence of an explicit agreement, communication or deliberate coordination between companies. However, collusion can happen automatically through algorithms, without the need for explicit agreement among actors (Ezrachi & Stucke, 2017), and current laws are not suitable to identify legitimate optimisation of algorithms and anti-competitive coordination among them. The European Commission and the US Federal Trade Commission have both highlighted algorithmic collusion as a space that needs to be regulated, but the means of enforcing such regulation are not yet developed (OECD, 2017).

One of the main problems lies with the intent. Traditional antitrust theory assumes "conscious human coordination. A reinforcement learning system, by contrast, doesn't need to be taught explicitly by the firms using the system to converge on a collusive equilibrium, and this makes it far more difficult to establish the intentionality that is traditionally required for anti-competitive conduct.

Explainability is another challenge related to this. Many AI pricing algorithms can be quite complex and not easily understandable. Such opacity creates informational asymmetries between firms and regulators and consumers, and Burrell (2016) asks whether less opaque AI systems can be developed to make it easier for regulators to determine if an algorithm has been designed to facilitate coordination.

The fast and high volume of trading makes algorithmic markets hard to monitor as well. AI systems continually and quickly adapt to changing circumstances, whereas regulatory investigations are slow. This can facilitate the existence of anti-competitive practices for a long time without them being discovered.

These problems have led to significant work in the scholarship literature concerning the need for major changes in antitrust law. The authors of the paper (Ezrachi and Stucke, 2017) contend that algorithmic collusion is not just a technologically new variant of an existing issue – it's a change in the structure of market coordination that demands a fundamental rethinking of competition policy beyond the traditional focus on human strategic conduct. On the regulatory side, some proposals aim at enhancing the regulatory framework through algorithmic auditing mechanisms (CMA, 2021) while others are focused on the transparency of AI pricing systems. However, over-regulation has its dangers, such as stifling innovation, or otherwise imposing undue burden on smaller companies versus larger, more-capable incumbents.

Digital platforms are also cross-border, and algorithmic pricing is often cross-border, with algorithms running across different jurisdictions with different competition laws. International cooperation and some harmonisation of approaches to AI regulation might be needed to implement effective governance in this area.

7. CRITICAL ANALYSIS: EMERGENT STRATEGIC COORDINATION

7.1 The Pro-Efficiency Case for Algorithmic Pricing

This is not to downplay the significant literature that views algorithmic price as welfare improving before turning to a critique of algorithmic coordination. While some economists use only the collusion risk metric to determine the value of dynamic pricing algorithms, they believe it serves the consumer and the market well to do so.

Algorithmic pricing is claimed to increase allocative efficiency by bringing prices more in line with the real-time supply and demand; for instance, Hall et al. (2015) show that, during periods of high demand Uber's algorithmic pricing results in lower waiting times for passengers and more drivers available. Second, algorithms are claimed to lower search and information expenses for consumers, because they allow them to compare prices in real time from a variety of platforms, offering a natural limit on continued price hikes. Third, by providing a more personalized price, dynamic pricing can lead to more consumer welfare gains, either through better price-setting to willingness to pay, or through a reduction in dead weight loss.

Brown and MacKay (2023) also claim that transparency can simultaneously limit the coordination power of the algorithm as well: if the price deviation can be seen on the spot, then a competitor is better able to buy at such a price and beat the collusion; in addition, it makes it easier for new entrants to compete on price in a transparent market. On this view, algorithmically managed markets may, in some circumstances, be more contestable than markets managed by human decision-makers.

These efficiency arguments deserve serious consideration, and the present review does not dispute that the gains they identify are real, and in many market structures may outweigh the risks associated with coordination. Rather, the argument developed in this section is that efficiency gains and anti-competitive risks can coexist within the same algorithmic systems, and that existing analytical and regulatory frameworks are not well equipped to distinguish between the two, or to address the latter without undermining the former.

7.2 Emergent Strategic Coordination: A Conceptual Framework

A significant limitation in the existing literature concerns how algorithmic collusion is conceptualised. Much of the literature analyses algorithmic coordination using concepts developed for traditional tacit collusion. This section argues that AI-based strategic behaviour constitutes a meaningfully different category of problem, referred to here as emergent strategic coordination: market coordination that is realised computationally through autonomous optimisation processes, rather than through intentional human decision-making.

This is not a claim that the problem of non-intentional algorithmic coordination has gone entirely unnoticed in the existing literature. Harrington (2018) develops an article-length treatment of competition law specifically directed at collusion by autonomous artificial agents. Ezrachi and Stucke (2016, 2017) distinguish between scenarios of algorithmic collusion according to the degree of human involvement and intent, an analysis set out at length in their 2016 book and developed further in their 2017 article. The

OECD (2017) and Deng (2018) likewise address coordination that arises without explicit human agreement. The contribution of this paper is therefore better understood as a synthesis and extension of these existing accounts, rather than as the identification of an entirely unaddressed problem. Specifically, this paper offers a structured three-dimensional articulation of the distinction between emergent and traditional coordination, extends the analysis to the role of platform architecture as an active factor in shaping coordination (rather than treating it as background context), and incorporates the more recent phenomenon of LLM-based pricing agents, which the earlier frameworks largely predate.

Three dimensions distinguish emergent strategic coordination from traditional tacit collusion as conventionally understood.

1. Intent and awareness: traditional accounts of tacit collusion assume that managers consciously recognise a shared interest in avoiding aggressive price competition and choose, without direct communication, not to compete on price. In emergent strategic coordination, by contrast, no individual within a firm need recognise that cooperative behaviour is occurring; the algorithm arrives at cooperative pricing because doing so maximises cumulative reward, and the humans deploying the algorithm may be entirely unaware that this dynamic has emerged.
2. Coordination mechanism: traditional tacit collusion relies on public, human-interpretable signals such as price leadership or industry-standard pricing norms. Gradient descent, reward shaping and pattern recognition over a large number of market observations, on the other hand, lead to emergent coordination. The outcome of the behaviour is not the result of any conscious strategic choice, but is rather embodied in learned weights and policy functions, meaning that this is a behaviour that is clearly different from those modes of signaling and parallel conduct that antitrust law was intended to uncover (Harrington, 2018).
3. Detectability and accountability: Usually, there is something that remains detectable and accountable about collusion, whether it is through communication, meetings, or pricing, when there is traditional collusion. Emergent coordination is much more difficult to observe, because it is a consequence of machine learning systems that are often not meaningfully auditable, even by the companies that put them into use (Burrell, 2016). The consequences of such a behavior are correspondingly spread out, and could be related to an algorithm developer, the company that implements the algorithm, the platform it runs on, or the provider of the underlying data, none of which would necessarily have an intentional misbehavior, as required by antitrust law.

These three factors may help to understand why, until now, the substantive evidence and computational models of algorithmic collusion have never led to established regulatory change or theory modifications: the current conceptual or legal framework was not developed with this phenomenon in mind.

This approach has implications for many aspects of platform design, as well as for designing various forms of recommendation systems, ranking algorithms for visibility, and data structure, all of which influence the conditions under which pricing algorithms learn their behavior. In this perspective, platform design is not a passive component of competitive interactions, but a central feature of the overall system that generates the coordination of algorithms, which can be neglected in most game-theoretic models that assume competing algorithms work in isolation from their platform context.

Last but not least, a small number of large technology companies are likely to exacerbate the risk of emergent strategic coordination because they will have larger data sets and superior data infrastructure, enabling them to implement better algorithms than their smaller rivals. Emergent strategic coordination might thus tend to consolidate current structural benefits in digital markets, rather than undermining them.

8. LIMITATIONS

A few important limitations to this review must be noted right off the bat. First, as a narrative, rather than systematic review, the choice of sources is based on the reviewer's sense of their relevance and importance and not necessarily a “reproducible search” with set inclusion and exclusion criteria, and so may differ slightly from the choice of a different reviewer who has access to the same broad literature. Any statements that can be made in this paper, whether relating to what the literature does or does not cover, should be understood in this context and as observations based on the literature referenced rather than as comprehensive statements about an entire field.

Second, much of the evidence on which the theory of algorithmic collusion is based is built on computational simulations, such as Calvano et al. (2020) and Klein (2021), rather than evidence from real-world markets. The conclusions of Werden and Froeb (2018) with respect to claims of increased concentration are more general: striking results of simulations or stylized settings do not necessarily imply that a phenomenon is a serious issue in real markets, which are more complex, more populated by competing firms, and subject to managerial oversight that simulations may not adequately represent. There is some real world evidence, from a single market (German retail petrol stations) and from a working paper that has not yet been peer-reviewed, that appears to be consistent with the simulation results, but it is unclear if the results would hold for other markets or sectors.

Third, the structure of emergent strategic coordination from Section 7 is put forward as an analytical and conceptual contribution, which is not meant to be an empirically tested theory. This paper does not and cannot answer the question as to whether the three dimensional distinction that it posits results, or does not, in materially different regulatory conclusions than the existing work by Harrington (2018) or Ezrachi and Stucke (2016, 2017).

Fourth, while regulatory features in this paper have been analyzed from a U.S. and EU perspective, very little attention is given to regulatory considerations elsewhere, including in fast digitalizing economies, which would warrant further examination in a separate paper.

9. FUTURE RESEARCH DIRECTIONS

Further work could usefully test the extent to which the dynamics identified in computational simulations (Calvano et al., 2020; Klein, 2021) are observed in real markets across a wider range of sectors than the single market examined by Assad et al. (2020). Research combining market-level data with variation in algorithmic adoption, in the manner of Assad et al. (2020), but extended to additional markets, would help to establish how far the simulation-based findings generalise.

Another key issue is the formal representation of platform architecture in models of algorithmic competition. Existing game-theoretic models tend to represent the different algorithms as competing in a relatively stable, isolated environment; future models could allow the recommendation systems, the feedback the data gives to the algorithm, and other choices made by platform designers to evolve, thereby changing the dynamics of their emergence.

Third, as algorithmic pricing is often also managed, adjusted, and even overridden to a certain degree in practice, studying how human management affects emergent coordination through algorithmic learning, either by interrupting it or reinforcing it, or by overcoming it, would help to deepen our understanding of how this kind of learning interacts with human management in authentic settings.

Fourth, given that pricing in digital markets is often entangled in advertising, logistics, and personalisation, the possibility of coordination spreading from one market to another related markets through the use of common data infrastructure deserves to be explored, as in the case of large digital conglomerates (Petit, 2020), or as it has been done in the context of multi-market competition.

Lastly, some regulators have been treated as rule-makers, in practice rather than as strategic actors, which means that future research could benefit from an evolutionary approach to the relationship between regulators and firms and algorithms, where regulators' interventions evolve over time based on the behaviour of firms and algorithms.

10. CONCLUSION

One of the biggest changes in competitive dynamics in today's market is the introduction of algorithmic pricing and AI-powered competitive strategies, where reinforcement learning systems, dynamic pricing algorithms, and platform ecosystems influence pricing decisions. While these technologies offer real advantages in terms of efficiency and responsiveness to market requirements, they also pose serious risks of tacit collusion, market concentration and regulatory failure.

This paper has covered the ways in which AI-based pricing systems are altering both the manner in which firms are competing in digital markets, and the speed at which these changes are occurring, all which is occurring at a much more rapid rate than the regulatory and theoretical "catch-up". In computational simulations, reinforcement learning agents have been demonstrated to reach a collusive equilibrium without communication and human intervention. As found in this review, there are existing documents demonstrating that this capacity exists (Calvano et al., 2020; Klein, 2021), but regulatory and theoretical frameworks have not advanced as quickly and the existing antitrust doctrine is still largely focused on human conduct.

This paper has covered the ways in which AI-based pricing systems are altering both the manner in which firms are competing in digital markets, and the speed at which these changes are occurring, all which is occurring at a much more rapid rate than the regulatory and theoretical "catch-up". In computational simulations, reinforcement learning agents have been demonstrated to reach a collusive equilibrium without communication and human intervention. As found in this review, there are existing documents demonstrating that this capacity exists (Calvano et al., 2020; Klein, 2021), but regulatory and theoretical frameworks have not advanced as quickly and the existing antitrust doctrine is still largely focused on human conduct.

The review has also pushed for the structural aspect – such as the design of digital platforms, data asymmetries, and visibility architecture – to be given more visibility in explaining the competitive outcomes, especially when compared to the existing literature, and to be further explored in future research.

The conclusions contained in this review are subject to serious restrictions (as discussed in Section 8): The method of the review is narrative and not systematic; The central body of evidence consists of computational simulations and does not rely upon any real marketplace data; The conceptual framework proposed here has not yet been tested empirically against the existing frameworks. Such restrictions highlight direction for future studies, rather than weaken or contradict the main claim of the paper.

Overall, the review indicates that algorithmic pricing is more than just a new technology applied to an old business model, and it necessitates new approaches to conceptualizing this kind of competitive market and new tools for regulating it, rather than tools built primarily for human decision makers.

REFERENCES

Assad, S., Clark, R., Ershov, D., & Xu, L. (2020). Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market. *SSRN Electronic Journal*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3682021

Autorité de la concurrence & Bundeskartellamt. (2019). *Algorithms and competition*. Paris and Bonn.

Belleflamme, P., & Peitz, M. (2021). *The economics of platforms: Concepts and strategy*. Cambridge University Press.

Brown, Z. Y., & MacKay, A. (2023). Competition in pricing algorithms. *American Economic Journal: Microeconomics*, *15*(2), 109–156.

Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society*, *3*(1), 1–12.

Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, *110*(10), 3267–3297.

Chen, L., Mislove, A., & Wilson, C. (2016). An empirical analysis of algorithmic pricing on Amazon Marketplace. In *Proceedings of the 25th International Conference on World Wide Web* (pp. 1339–1349). International World Wide Web Conferences Steering Committee.

CMA. (2021). *Algorithms: How they can reduce competition and harm consumers*. Competition and Markets Authority.

Deng, A. (2018). What do we know about algorithmic tacit collusion? *Antitrust Chronicle*, *2*(1), 88–99.

Ezrachi, A., & Stucke, M. E. (2016). *Virtual competition: The promise and perils of the algorithm-driven economy*. Harvard University Press.

Ezrachi, A., & Stucke, M. E. (2017). Artificial intelligence & collusion: When computers inhibit competition. *University of Illinois Law Review*, *2017*(5), 1775–1810.

Fish, S., Gonczarowski, Y. A., & Shorrer, R. I. (2024). *Algorithmic collusion by large language models*(arXiv:2404.00806). arXiv. https://arxiv.org/abs/2404.00806

Fudenberg, D., & Tirole, J. (1991). *Game theory*. MIT Press.

Hall, J. V., Horton, J. J., & Knoepfle, D. T. (2015). *Pricing in designed markets: The case of ride-sharing* [Unpublished manuscript]. Uber Technologies.

Harrington, J. E. (2018). Developing competition law for collusion by autonomous artificial agents. *Journal of Competition Law & Economics*, *14*(3), 331–363.

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

Lepore, J. J. (2021). *Algorithmic collusion in a volatile market: Q-learning under incomplete information* [Working paper].

Miklos-Thal, J., & Tucker, C. (2024). *AI, algorithmic pricing, and collusion* (NBER Working Paper Series). National Bureau of Economic Research.

OECD. (2017). *Algorithms and collusion: Competition policy in the digital age*. OECD Publishing.

Parker, G., Van Alstyne, M., & Choudary, S. P. (2016). *Platform revolution*. W. W. Norton.

Petit, N. (2020). *Big tech and the digital economy*. Oxford University Press.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.

Tirole, J. (2015). Market failures and public policy. *American Economic Review*, *105*(6), 1665–1682.

Varian, H. R. (2019). Artificial intelligence, economics, and industrial organization. *Innovation Policy and the Economy*, *19*(1), 399–419.

Werden, G. J., & Froeb, L. M. (2018). Don't panic: A guide to claims of increasing concentration. *Antitrust Magazine*, *33*(1), 74–79.