CIRANO

Knowledge into action

DECIPHERING ALGORITHMIC COLLUSION: INSIGHTS FROM BANDIT ALGORITHMS AND IMPLICATIONS FOR ANTITRUST ENFORCEMENT

FRÉDÉRIC MARTY
THIERRY WARIN

CS

2023s-26
WORKING PAPER

ISSN 2292-0838 (online version)

Deciphering Algorithmic Collusion: Insights from Bandit Algorithms and Implications for Antitrust Enforcement

Frédéric Marty [*], *Thierry Warin*[†]

Abstract/Résumé

This paper examines algorithmic collusion from legal and economic perspectives, highlighting the growing role of algorithms in digital markets and their potential for anti-competitive behavior. Using bandit algorithms as a model, traditionally applied in uncertain decision-making contexts, we illuminate the dynamics of implicit collusion without overt communication. Legally, the challenge is discerning and classifying these algorithmic signals, especially as unilateral communications. Economically, distinguishing between rational pricing and collusive patterns becomes intricate with algorithm-driven decisions. The paper emphasizes the imperative for competition authorities to identify unusual market behaviors, hinting at shifting the burden of proof to firms with algorithmic pricing. Balancing algorithmic transparency and collusion prevention is crucial. While regulations might address these concerns, they could hinder algorithmic development. As this form of collusion becomes central in antitrust, understanding through models like bandit algorithms is vital, since these last ones may converge faster towards an anticompetitive equilibrium.

Cet article examine la collusion algorithmique d'un point de vue juridique et économique, en soulignant le rôle croissant des algorithmes dans les marchés numériques et leur potentiel de comportement anticoncurrentiel. En utilisant comme modèle les algorithmes de bandits, traditionnellement appliqués dans des contextes de prise de décision incertaine, nous mettons en lumière la dynamique de la collusion implicite sans communication manifeste. Sur le plan juridique, le défi consiste à discerner et à classer ces signaux algorithmiques, notamment en tant que communications unilatérales. D'un point de vue économique, la distinction entre une tarification rationnelle et des schémas de collusion devient complexe dans le cas de décisions pilotées par des algorithmes. L'article souligne qu'il est impératif que les autorités de la concurrence identifient les comportements inhabituels sur le marché, en suggérant de transférer la charge de la preuve aux entreprises qui pratiquent la tarification algorithmique. Il est essentiel de trouver un équilibre entre la transparence des algorithmes et la prévention de la collusion. Si les réglementations peuvent répondre à ces préoccupations, elles risquent d'entraver le développement des algorithmes. Cette forme de connivence devenant un élément central de la lutte antitrust, il est essentiel de comprendre les modèles tels que les algorithmes de bandits, car ces derniers peuvent converger plus rapidement vers un équilibre anticoncurrentiel.

---

[*]CNRS – GREDEG – Université Côte d'Azur - CIRANO, Montréal
[†]HEC Montréal - CIRANO, Montréal

Keywords/Mots-clés: Algorithmic Collusion, Bandit Algorithms, Antitrust Enforcement, Unilateral Signals, Pricing Strategies / Collusion algorithmique, algorithmes de bandits, application de la législation antitrust, signaux unilatéraux, stratégies de fixation des prix

JEL Codes/Codes JEL: L13, L41, K21

Pour citer ce document / To quote this document

Marty, F., & Warin, T. (2023). Deciphering Algorithmic Collusion: Insights from Bandit Algorithms and Implications for Antitrust Enforcement (2023s-26, Cahiers scientifiques, CIRANO.) https://doi.org/10.54932/IWPG7510

I – Introduction

Algorithmic collusion, characterized by the confluence of traditional collusion dynamics and the innovative dimension of artificial intelligence, has emerged as a pressing concern in contemporary competition landscapes (Ezrachi and Stucke, 2016). While the literature has extensively addressed this phenomenon, the specific dimension of coordination through algorithm-produced signals remains relatively underexplored. Our research primarily concentrating on the assessment of the capacity of algorithms to produce signals facilitating such collusive equilibria in oligopolistic markets and the challenges raised for antitrust laws enforcers.

Historically, the Air Traffic Publishing Co case in 1994 stands as a seminal instance of such algorithmic collusion, where airlines utilized a digital platform to unilaterally broadcast prospective price augmentations, leading to an equilibrium achieved via a cascade of price signals. This exemplifies the potential of algorithms in engendering tacit collusion. Even if online prices are intrinsically more differentiated and dynamic than offline ones, collusive equilibria may also be reinforced by the digital age's facilitation of real-time price experimentation and competitor monitoring. This agility in pricing may also introduce complexities for competition authorities in overseeing price alterations, a challenge mirrored in high-frequency trading where signals, though cryptic to regulators, are discernible to market players (Arena et al. 2018). The paradox of the current situation lies in the fact that it is possible to reconcile a high degree of price differentiation and variability with the firms' ability to identify distinct signals that allow for the identification of a focal point for possible coordination.

The history of antitrust caselaw has grappled with analogous challenges. “The Gary Dinners” witnessed oligopolistic entities disseminating future strategies through unilateral declarations, culminating in a landmark Supreme Court ruling in 1920 (Page, 2009). More contemporaneously, the European Commission's 2016 adjudication of the Container Shipping case reaffirmed the potential repercussions for firms abetting collusion via unilateral communication.

This brings us to an important aspect of antitrust decision-making practices, where unilateral acts can serve as facilitators for achieving collusive equilibria. These can manifest through acts of unilateral communication that reveal a company's future strategy, or through unilateral transparency that enables competitors to anticipate and align with this strategy.

In the intricate domain of pricing algorithms, Salcedo (2015) postulated a model of algorithmic collusion, suggesting that firms might intentionally render their algorithms transparent to rivals. This strategic transparency enables competitors to recalibrate their algorithms and pricing strategies, fostering parallelism. Alternatively, subtle signals might act as overtures for coordination, progressively gravitating towards a collusive equilibrium. In finance, this idea is often explored in the context of market microstructure and high-frequency trading, where traders adjust their bids and asks based on the ongoing flow of market information and the actions of other market participants. Models in this domain frequently analyze how equilibrium prices are reached through a process of continuous adjustment, reflecting the dynamic and interactive nature of financial markets (Biais, Hombert, et Weill 2014).

While signaling strategies are not exclusive to the use of algorithms, it's important to note that these strategies can be significantly 'enhanced' or 'augmented' by the capabilities inherent in algorithmic systems.

Our research delves into the bandit algorithm's implications in this context. The bandit algorithm, rooted in probability theory and statistics, addresses the exploration-exploitation dilemma, determining optimal strategies under uncertainty. In the realm of algorithmic collusion, the bandit approach can elucidate how firms might balance between exploring new pricing strategies and exploiting known beneficial ones, especially when navigating the landscape of unilateral signals.

The specificity of bandit algorithms compared to self-reinforcing machine learning lies in their unique design to optimize decision-making under uncertainty, balancing exploration and exploitation. This is a departure from self-reinforcing ML algorithms that predominantly learn from and reinforce patterns in historical data. Bandit algorithms, as exemplified in the work of Agrawal and Goyal (2012), excel in environments where they must make decisions with minimal prior data, adapting continuously based on immediate feedback. This adaptability, however, introduces significant collusive risks and supervisory challenges. As already demonstrated by Waltman et al. (2008) and later by Calvano et al. (2020) about Q-learning, algorithms can inadvertently learn collusive patterns, presenting difficulties for competition authorities to discern competitive behaviors from collusive ones due to their dynamic nature.

It is technically plausible that bandit algorithms can converge faster than self-reinforcing machine learning models in certain contexts. This can be attributed to the fundamental operational differences between these two types of algorithms. Bandit algorithms are designed for optimal decision-making in environments with uncertain and variable outcomes. Their strength lies in their ability to quickly adapt to new information, adjusting their strategies based on immediate feedback. This adaptability is particularly evident in environments where real-time decision-making is crucial and where available data may be sparse or non-stationary.

In contrast, self-reinforcing machine learning models typically rely on large datasets and often seek to identify and reinforce patterns based on historical data. While powerful in scenarios with abundant data and well-defined patterns, these models might require more time to converge in environments where the data is continually changing or where immediate feedback is integral to the decision-making process.

Numerical simulations can indeed be used to demonstrate this difference in convergence rates. By simulating environments where conditions rapidly change or where only limited data is available, one can observe how bandit algorithms adjust their strategies more swiftly to optimize outcomes. In contrast, self-reinforcing ML models might take longer to identify and adapt to these changes due to their reliance on historical data patterns.

This faster convergence of bandit algorithms can be particularly relevant in dynamic market environments, such as digital markets, where conditions frequently change and immediate decision-making is crucial. However, it is important to note that the speed of convergence does not necessarily imply superiority in all aspects. While bandit algorithms may excel in rapidly changing environments, self-reinforcing ML models might be more effective in stable environments with vast historical data.

The overall contribution and positioning of the paper revolve around demonstrating how bandit algorithms can amplify existing risks of collusion in digital markets. The paper delves into the formal model underlying these algorithms to illustrate how such risks are heightened compared to standard ML algorithms. It argues that bandit algorithms not only perform better in real-time decision-making but also present greater challenges in interpretability and regulatory oversight, making them more difficult to decipher for competition authorities and potentially exacerbating undetected collusion risks.

In terms of positioning in relation to existing literature, the paper distinguishes itself by investigating whether other studies have addressed similar questions using bandit algorithms, particularly in the context of antitrust implications. Additionally, examining literature from other fields, such as finance, where bandit algorithms are applied for signal detection, provides further insights. This cross-disciplinary review helps in comprehending the functioning of bandit algorithms in different contexts, shedding light on potential implications for collusive risks and regulatory challenges in digital markets.

This paper contributes significantly by exploring the unique characteristics of bandit algorithms, their potential to heighten collusive risks in digital markets, and the challenges they pose for regulatory oversight, thus addressing a critical gap in the existing academic discourse.

Subsequent sections will delve into the intricacies of how unilateral collusive contract offers might be orchestrated in oligopolistic settings and the potential countermeasures competition regulators might deploy. Section 2 suggests that the ability to achieve collusive equilibria can be enhanced by the use of bandit algorithms. Section 3 outlines the challenges in terms of enforcing competition rules to characterise these trajectories as anti-competitive and to sanction them. Section 4 discusses the economic challenges raised by such a potential dynamic.

II – Augmented collusive capacities through algorithms? Investigating the case of bandit-algorithms.

A- Algorithmic collusion through signals: the state of play

The challenges posed by tacit collusion, especially when algorithms are involved, are continually evolving. Ivaldi et al. (2007) highlighted the inherent instability of these collusive equilibria, which complicates their maintenance for oligopolists and simultaneously poses challenges for competition authorities in their identification and regulation. Furthermore, the legal prerequisites for establishing anti-competitive intent during litigation add another layer of complexity. As De Coninck (2016) noted, competition authorities often target the underlying practices that facilitate such collusion.

Recent contributions, such as those by Xu, Lee, et Tan (2023), and Dorner (2021), have provided deeper insights into the dynamics of algorithmic collusion. They have emphasized the role of recommendation algorithms, the potential of self-learning algorithms in real markets, and the credible threat posed by such collusion, respectively.

The literature on tacit collusion, particularly when mediated by algorithms, is vast and multifaceted. A significant portion of this literature has delved into the role of complex machine learning algorithms in facilitating such collusion. Notably, Ezrachi and Stucke (2017), Mehra (2016), and Calvano and al. (2019) have posited that self-learning algorithms can spontaneously converge to a tacit collusive equilibrium, even if their initial parameters differ. This convergence occurs without any direct or indirect communication between the firms, thereby challenging traditional methods of detecting and prosecuting anti-competitive behavior.

Recent studies, such as those by Harrington (2018), Schwalbe (2018), Calvano et al. (2019, 2020), and Assad et al., (2023), have further emphasized the competitive risks posed by pricing algorithms. The consensus is growing around the potential of these algorithms to autonomously reach a tacit collusive equilibrium without explicit communication. This blurs the distinction between parallel behavior and concerted practices (Calvano et al. 2020).

While the aforementioned studies have focused on the emergent properties of complex machine learning systems, this paper seeks to shed light on a less technologically advanced yet potentially equally effective form of algorithmic collusion. Here, algorithms are employed not for their self-learning capabilities but for their ability to transmit signals between market operators. This signaling allows these operators to identify focal points for coordination, potentially enabling firms to reach a collusive equilibrium and thereby potentially undermining market competition.

Such an anti-competitive strategy has been observed in some oligopolistic markets, in which companies wish to 'signal' their future strategies to their competitors to facilitate the identification of a focal point in coordination (Brown, Eckert, and Silveira 2023). Such a practice was detected in Canada in 2013 in the province of Alberta on the wholesale 'day-ahead' electricity market. On electricity markets, generating unit operators announce hour-by-hour price/capacity pairs for the following day, which are used by the grid operator to establish an order of economic priority. These announcements are, however, conditional, as operators can revise their offers (to take account of production contingencies) up to two hours beforehand. As bids were published anonymously by the operator, the strategy of some firms was to allow their bids to be identified indirectly, to 'propose' a cooperative solution to their competitors. In the absence of a countersignal, the firm could abandon its proposal, which therefore had no possible impact on price.

One way of signaling a bid, for example, might be to use specific (or mathematically improbable) frequencies of digits after the decimal point. Collusion strategies in the Libor cartel case (Abrantes-Metz et al., 2011), on the Nasdaq (Christie and Schultz, 1994) or on the price of aluminum on the London Metal Exchange (Sama, 2014) used such signaling tactics.

This method of collusion does not hinge on the intricacies of machine learning or autonomous learning processes. Instead, it relies on a relatively straightforward mechanism of signal transmission. Such a perspective underscores the idea that even simple algorithms, if used strategically, can facilitate anti-competitive behavior.

The integration of algorithms in market collusion presents a complex landscape. Algorithms, particularly those akin to bandit algorithms, can serve as communication conduits, emitting signals discernible to competitors but potentially elusive to competition authorities. The pivotal question centers on the ability of algorithmic signals to act as unilateral collusion invitations and acceptance indicators, all while circumventing the appearance of mutual agreement information exchange. Moreover, the capacity of algorithms, especially those based on the multi-armed bandit framework, to amplify signal clarity and filter out disruptive "noise" is of paramount importance.

Collusion can either emerge from sheer unilateral transparency without explicit information exchange or from bilateral transparency where information is reciprocally shared. Such maneuvers, from a competition law lens, might be construed as anti-competitive. Introducing artificial transparency might be perceived as a collusion intent signal, especially in parallel behavior contexts.

In terms of collusion, several scenarios can be envisioned. One involves a professional association mitigating uncertainty regarding competitors' future actions. Another sees a benchmark firm signaling industry median firm preferences. A third scenario involves a Stackelberg leader, a dominant entity seeking alignment with peripheral competitive fringe firms. These fringe firms, as Gal (2019) illustrates, can adopt cooperative stances via pricing algorithms. A Stackelberg leader might employ an advanced pricing algorithm, while competitors opt for simpler, leader-adaptive algorithms using identical databases, expediting convergence.

In this milieu, algorithms, especially those resembling the bandit algorithms, can extend cooperation invitations to competitors. This doesn't necessarily culminate in formal cartel formation. Instead, the algorithmic signals aim to pinpoint coordination focal points through iterative processes, hinging on competitor feedback. This can be analogized to a unilateral contract offer, with competitors potentially reciprocating by adjusting prices. Absent such reciprocation, the initiating firm might foresee a unilateral price hike resulting in sales plummeting.

Economically, this can be dissected into two scenarios. In one, a Stackelberg leader signals future prices, indicating that undercutting would be perceived as a price war provocation. This strategy hinges on aggressive behavior threats. In another, the leader proposes a universally profitable price, emphasizing collective over individual profit. The leader's strategic intent dictates the scenario choice.

The leader can either announce immediate prices or hint at future adjustments. Immediate changes should occur during low-sales periods, allowing competitor algorithms to swiftly detect and adapt. Algorithmic price adjustments streamline this process. Alternatively, price shifts can be postponed, with unilateral price hikes contingent on unanimous future price direction changes by all firms. The Airline Tariff Publishing Co. case exemplifies such dynamics.

Tracing back to the Airline Tariff Publishing Co. case, which culminated in a 1994 settlement, the company disseminated extensive future pricing data. This data richness facilitated indirect inter-company communication (Borenstein, Kwoka, et White 2004), prefiguring algorithmic collusion investigations.

However, volatile price movements can introduce noise, complicating signal interpretation. Algorithmic pricing can lead to complex price dynamics and the difficulty for both competitors and regulators in discerning whether a particular pricing pattern is the result of competitive behavior or a signal of collusion (Calvano et al. 2019).

B- Bandit-Algorithms and Collusive Equilibria

In the realm of competition law, understanding the role of algorithms in fostering and maintaining collusive equilibria is crucial. Collusion can manifest through explicit signals that foster mutual understanding among competitors, or through AI algorithms that spur spontaneous tacit collusion. Cooper et al. (1989) posited that unilateral communication, particularly when aligned with collective interests, simplifies the path to collusive equilibrium compared to bilateral communication, thereby enhancing the likelihood of mutually beneficial outcomes.

The iterative nature of such interactions strengthens the credibility of unilateral declarations by a leading entity. If the leader's commitments don't align with subsequent actions or send misleading signals, their reputation could suffer, potentially undermining future credibility. This ongoing interaction builds trust, and the willingness of a leader to take unilateral risks underscores their commitment to cooperation, as explored by Salcedo (2015).

In the context of signaling models, algorithmic pricing can foster collusive tendencies by creating a form of artificial transparency. This is notable in the case of pricing algorithms based on self-reinforcing machine learning. While parallel algorithms may facilitate collusive equilibria through tit-for-tat strategies as observed by Axelrod and Hamilton (1981), self-learning algorithms might offer even more advantageous outcomes. This aspect is elaborated by Hingston and Kendall (2004), who demonstrate how these algorithms can adapt and optimize strategies for competitive advantage.

Crandall et al. (2018) empirically showed that autonomously learning algorithms could develop enduring cooperative relationships with both machines and humans. However, the debate about the genuine cooperative capacity of AI algorithms persists. Notably, tacit collusion often relies on intuition, cultural norms, and the ability to decipher subtle signals from competitors. In this setting, non-binding signals, or "cheap talks," become crucial.

Yet, algorithm-driven collusion introduces significant complexities in competition law. Strategies rooted in creating unilateral transparency for competitors, without direct exchanges or overt agreements, challenge traditional anti-competitive paradigms. This aspect demands careful legal scrutiny.

In essence, algorithms, regardless of their nature, have the potential to catalyze the formation of collusive equilibria. Algorithms that signal, in particular, are intriguing from a competition law perspective, raising questions about strategies rooted in unilateral transparency without overt exchanges. Thorough research is essential to unravel the intricacies of algorithmic collusion in the context of competition law.

Shifting focus to the bandit algorithm, this technique, often associated with multi-armed bandit problems, addresses decision-making amid uncertainty. Originating from a hypothetical scenario where one must choose the most rewarding slot machine arm to pull, the bandit algorithm confronts the exploration-exploitation dilemma: whether to explore all options to understand their reward distribution or exploit the most rewarding option based on current knowledge. This dilemma is addressed in various strategies like ε-greedy, Upper Confidence Bound (UCB), and Thompson Sampling, as detailed by Sutton and Barto (2018).

Bandit algorithms find diverse applications, from guiding patient treatments in healthcare to managing resource allocation in web services. In online advertising, as Li et al. (2010) highlight, these algorithms determine which ads to display to maximize click-through rates. In algorithmic pricing, a domain where bandit algorithms particularly excel, each pricing strategy is an 'arm', with the 'reward' being the resultant profit. Babaioff et al. (2015) demonstrate how these algorithms navigate consumer response uncertainties, balancing exploration of new pricing strategies with exploitation of known profitable ones.

In the formal model of the multi-armed bandit problem, an agent interacts with an environment over time, choosing actions to maximize cumulative rewards. The goal is to minimize regret, defined as the difference between potential and actual rewards. Various algorithms, including ε-greedy, UCB, and Thompson Sampling, tackle this problem by balancing exploration and exploitation. UCB1, for instance, is a heuristic approach that achieves logarithmic total regret, indicative of its efficiency, as shown by Auer, Cesa-Bianchi, and Fischer (2002).

Bandit algorithms differ significantly from traditional ML algorithms in real-time adaptability, efficiency in balancing exploration and exploitation, and less reliance on extensive datasets. These attributes, explored by researchers like Agrawal and Goyal (2012) and Bubeck and Cesa-Bianchi (2012), make them uniquely suited for dynamic environments. However, their evolving nature also complicates the interpretation of their strategies, especially in regulatory and antitrust contexts, as discussed by Harrington (2018). This complexity necessitates advanced monitoring techniques to effectively detect collusion and interpret the underlying intentions of these algorithms.

application within the field of reinforcement learning. Q-learning, a model-free reinforcement learning algorithm, is designed to determine the best action to take in a given state. It functions without a model of the environment, hence its designation as 'model-free', and it is capable of handling environments with stochastic transitions and rewards. The principle behind Q-learning is the estimation of the value of a state-action pair, or Q-value, which reflects the expected utility of taking a specific action in a given state and then adhering to the optimal policy thereafter (Watkins & Dayan, 1992). This methodology enables Q-learning to effectively address complex decision-making tasks in environments with multiple states and actions.

Bandit algorithms, in contrast, are more straightforward and are primarily employed in multi-armed bandit problems. These problems involve an agent selecting from several choices, each with unknown and potentially varied reward distributions. The core challenge in these scenarios is balancing exploration (experimenting with different choices to discover their rewards) and exploitation (opting for the choice that has historically offered the best rewards). Bandit algorithms are typically used in simpler, stateless contexts, making them distinct from Q-learning, which factors in the state of the environment in its decision-making process. Their applications include areas like A/B testing, adaptive routing, and personalized recommendations, where the primary objective is to maximize rewards in situations of uncertainty (Auer, Cesa-Bianchi, & Fischer, 2002).

The contrast between Q-learning and bandit algorithms can be seen in several dimensions. Q-learning is more complex and appropriate for state-dependent decision-making environments, while bandit algorithms are simpler and designed for stateless decision-making contexts. Q-learning algorithms consider the state of the environment when making decisions, whereas bandit algorithms do not. Both approaches address the exploration-exploitation trade-off, but they do so in different contexts and with varying degrees of complexity. Q-learning is part of a broader suite of reinforcement learning algorithms and is suited for more complex tasks in structured environments. Bandit algorithms, on the other hand, focus on simpler reward maximization problems.

C. Practical Implications and Algorithmic Coordination

Offline unilateral price announcements have inherent legal risks and may necessitate iterative adjustments. Excessive transparency might also compromise coordination efficacy. To navigate these challenges, algorithmic price monitoring and adjustment tools, especially those resembling bandit algorithms, can bolster coordination efficiency. They facilitate real-time price monitoring and rapid market signal-based adjustments.

However, the legal ramifications of algorithmic coordination are non-trivial. Unilateral algorithmic communication must steer clear of overt collusion or anti-competitive behavior contravening competition laws. Competition authorities must meticulously scrutinize market player signals and actions to ascertain potential anti-competitive conduct.

Bandit algorithms are particularly suited to the study of algorithmic collusion due to their capacity to navigate uncertainty and limited information, hallmarks of market environments. These algorithms adeptly balance the exploration of new strategies with the exploitation of known profitable ones, mirroring the decision-making process firms undergo when setting prices amidst opaque competitor actions and market responses.

In the realm of adaptive strategies, bandit algorithms stand out for their ability to iteratively learn which actions yield the highest rewards. This attribute is crucial in market settings, allowing pricing algorithms to evolve towards profit-maximizing behaviors that could inadvertently result in tacit collusion, all without any explicit coordination between competing entities.

The potential of bandit algorithms to decode signals from market competitors is particularly pertinent. They are capable of interpreting market fluctuations as indicators of the strategic shifts of rivals, enabling them to adjust their pricing strategies in a responsive manner. This capability is central to understanding how firms might employ subtle pricing cues to align on more profitable outcomes without direct communication, a concern for regulatory bodies.

Moreover, bandit algorithms can systematically explore and potentially exploit collusive equilibria. Through a process of trial and error, they can identify pricing strategies that, when mirrored by competitors, lead to mutually beneficial, albeit collusive, outcomes. This exploration is key to understanding the pathways through which algorithms may autonomously converge on anti-competitive practices.

The dynamic nature of markets, characterized by shifting consumer preferences and varying external conditions, is where bandit algorithms truly excel. They are designed to adapt in real-time, constantly refining strategies to align with the current state of the market, ensuring efficiency and responsiveness in pricing decisions.

The implications of bandit algorithms extend into policymaking. A thorough grasp of how these algorithms function and the conditions under which they may lead to collusion is invaluable for regulators. Such knowledge is essential for developing measures that can preempt and mitigate the risk of anti-competitive behavior.

Lastly, the deployment of bandit algorithms in setting prices also raises ethical questions regarding the responsibility of firms to prevent consumer harm. Insights gleaned from the study of these algorithms can contribute to the establishment of ethical standards for the application of artificial intelligence in commercial settings.

Bandit algorithms offer a robust framework for probing the emergence and dynamics of algorithmic collusion. Their inherent learning and adaptive capabilities provide a lens through which one can examine the potential for autonomous systems to engage in market behaviors that undermine competitive integrity. Understanding these mechanisms is crucial for shaping business strategies and regulatory frameworks that uphold the principles of market competition.

So, while algorithms, especially those akin to bandit algorithms, can potentiate coordination and cooperative equilibrium, signal interpretation challenges and competition law adherence are paramount. Algorithmic tools can augment coordination efficiency, but their deployment must respect legal and ethical boundaries to preserve market competition integrity.

III. Addressing signaling collusions risks: assessing the legal challenges.

From a legal standpoint, there have been instances where algorithms have been implicated in collusion, as highlighted by the OECD (2023). These range from upholding cartel agreements to orchestrating collusive agreements via hub & spoke schemes. The Air Tariffs Publishing Co. case serves as a pertinent example of the potential misuse of algorithms in guiding collusive behavior. Another dimension to consider is the potential of algorithms to facilitate collusion either by interfacing with rival algorithms or by emitting signals that establish a collusion's focal point, as seen in the ATPCO case.

The intersection of competition law and economic analysis is pivotal when addressing collusive signals.

From a legal perspective, discerning and evaluating collusive signals presents formidable challenges. Algorithm-driven collusion might encompass unilateral actions fostering transparency amongst competitors without manifest evidence of explicit agreements. This engenders queries regarding the categorization of such behavior and its alignment with anti-competitive practices (Harrington 2018).

Economically, it's imperative to scrutinize the incentives and strategic maneuvers of involved entities. Economic frameworks can elucidate the pros and cons of signal-facilitated collusion. Such models can probe the ramifications of signaling algorithms on market outcomes, pricing paradigms, and consumer welfare. Furthermore, understanding the conditions fostering effective collusive signals—such as signal credibility, recurrent interactions, and market transparency—is essential. Grasping the nuances of signal-based collusion necessitates an evaluation of algorithms' influence on market dynamics and the repercussions for market efficacy and competition.

Moreover, assessing the repercussions of collusive signals on market competition and consumer welfare is paramount. While collusion might culminate in escalated prices or diminished product quality, gauging the holistic welfare effects and potential competition detriments is crucial. Economic evaluations can quantify the probable economic detriment instigated by collusion, thereby informing the formulation of robust competition policies and enforcement strategies (Schlechtinger et al. 2023). algorithms can enter a collusive state and charge supra-competitive prices without explicit communication.

Unilateral communication is pivotal for coordination focal point identification. The concept of signaling future prices can indeed be construed as a unilateral offer to enter into a collusive agreement, which is antithetical to the principles of competitive markets. In this framework, such signaling is theorized to serve as a tacit communication strategy among competitors, effectively circumventing the explicit verbal or written contracts that competition law explicitly prohibits. This interpretation aligns with a broader understanding of market signaling and tacit collusion in economic theory. Tacit collusion, unlike explicit collusion, does not involve direct communication between parties to set prices or output. Instead, it occurs when businesses engage in practices that lead to a mutual understanding of shared market behavior without explicit agreement, which may be achieved through price signaling (Hansen, Misra, et Pai 2021). Posner's assertion that such signals constitute a unilateral contract offer is grounded in the tenets of contract law, where a unilateral contract is formed when one party makes a promise in exchange for the performance of an act by another party. In the context of price signaling, the 'promise' is the suggested future price, and the 'performance' is the adherence to this price by competitors, thereby creating a non-verbal collusive environment. This perspective is of considerable importance in the application of antitrust laws, as it frames certain market behaviors in terms of contract law, thereby potentially expanding the scope of actions considered collusive and thus anticompetitive. Antitrust authorities may, therefore, scrutinize such behaviors closely, despite the lack of explicit communication between firms. Understanding the subtleties of such market dynamics is crucial for both regulators and businesses to navigate the legal landscape effectively.

The European Commission posits that even unilateral signals might be construed as part of a collusive strategy. The Container Shipping decision delineates that any form of contact between operators intending to sway market behavior or divulge future conduct is strictly proscribed (§33). Article 101 penalizes collusive endeavors that, while not necessitating explicit agreements, entail the conscious adoption of collusive mechanisms that streamline commercial behavior coordination (§34).

Information exchanges are deemed anti-competitive if they artificially diminish the profound uncertainty underpinning firm operations, thereby amplifying tacit collusion risks. Particularly when information pertains to future strategies of competitors, its collusive potential escalates. EU competition law doesn't confine concerted practices to solely distortion-prone information; firms unable to substantiate their non-reliance on such information might be held accountable (§36).

In the Container Shipping case, individual firm announcements are perceived as communicative gestures, probing competitors' receptiveness to prospective behavioral shifts. While competitor responses might be indirect, their unilateral proclamations regarding future actions converge, facilitating cooperative adjustments. Thus, a firm's unilateral communication might be deemed a concerted practice as per the European Commission's guidelines on horizontal agreements, especially when one entity divulges strategic data to competitors who assimilate it (European Commission 2011).

The crux of concerted practice hinges on actions that artificially curtail strategic uncertainty for rival firms. Any direct or indirect contact that sways market behavior or unveils intended or prospective conduct, thereby enabling collusive outcomes, falls under this purview (§61). The ostensibly unilateral nature of such communication must be evaluated in the context of its repercussions on the future strategies of signal-receiving firms. A firm that procures such data from a competitor, even if unsolicited, is presumed to have assimilated the information and recalibrated its market behavior unless explicitly stated otherwise (§62).

Even when communication doesn't entail direct exchanges but transpires via press releases, as in the Container Shipping case, the Guidelines account for concerted practices. While genuinely public unilateral announcements, like those in newspapers, typically don't constitute concerted practices under Article 101(1), under certain circumstances, the prospect of identifying a concerted practice remains. For instance, if subsequent public proclamations by other competitors ensue the initial announcement and encompass strategic reactions and recalibrations, it might hint at a strategy geared towards achieving mutual coordination understanding (European Commission 2011).

Thus, the European Commission's stance perceives unilateral signals as potential collusive strategy components. Information exchanges, irrespective of their modality, if they curtail uncertainty and bolster coordination amongst competitors, might be subjected to scrutiny as a concerted practice under Article 101 of EU competition law.

Algorithmic price modifications, especially those entailing iterative adjustments in response to prior signals, might be categorized as collusive endeavors. However, it's pivotal to recognize that monitoring competitor prices and adjusting prices, provided they align with a logical market strategy, aren't intrinsically anti-competitive. Economic entities reserve the right to astutely recalibrate their prices based on the extant and anticipated actions of their competitors, as articulated by the European Commission (2017b).

Numerous cases of collusion using algorithms have already been observed, notably through the production of signals (Marty and Warin, 2023). For instance, in a case involving Lithuanian travel agencies, the stability of the collusive equilibrium was consolidated by the software system sending out a signal (in the form of e-mails) as soon as the operator validated a discount above a given threshold (EU Court of Justice, Eturas, 21 December 2016, Judgment, C-74/14).

As the EU Court of Justice stated: “The investigation carried out by the Competition Council established that, as a result of the technical modifications made to the E-TURAS system following the dispatch of the message at issue in the main proceedings, although the travel agencies concerned were not prevented from granting their customers discounts greater than 3%, they were nevertheless required to take additional technical steps in order to do so” (§12).

Similarly, in the case of the collusive signals produced on the Alberta wholesale electricity market cited above (Brown et al., 2023), the electricity regulator was able to identify tagging patterns whereby powerful market players indirectly revealed their identity to competitors in order to announce their intention to maintain high prices the following day.

The competitive treatment of algorithmic pricing signals remains nebulous. While it's ambiguous whether advanced tools that monitor and infer from another firm's prices would qualify as "communication" under Article 101, the potential for innovative and unconventional interactions meeting the "communication" criteria cannot be entirely dismissed (European Commission, 2017a, §33).

In collusion contexts, a mutual intent to disseminate information doesn't necessitate reciprocity, as underscored by the Conseil de la concurrence (2007). A unilateral information provision suffices, with the intent of the mutual understanding confined to information sharing. In cartel agreement scenarios, the coordination objective is to eradicate uncertainty regarding competitor behavior, achievable via information exchanges or unilateral market participant actions.

Under French competition law, a firm can be penalized for conveying estimates to competitors within a public contract context, as it contributes to uncertainty elimination and establishes a coordination focal point. Reciprocity between firms isn't a prerequisite for characterizing an information exchange in EU law. The concerted practices concept doesn't presuppose reciprocal interactions. Instead, the criterion is met when one competitor divulges its future market intentions or actions to another, even if it's a unilateral information transmission via a private channel. This added collusion dimension renders it arduous for the receiving firm to evade liability, as it consents to procure information that diminishes uncertainty in a competitive scenario (European Commission 2011).

Thus, while algorithmic price modifications might be subjected to competition law scrutiny, treating such signals as collusive practices is contingent on various factors, including market strategy rationality, communication nature, and uncertainty reduction amongst competitors. Reciprocity isn't invariably mandated, and unilateral information transmission might be deemed a collusion component.

It's quintessential to differentiate between parallel behavior and anti-competitive intent, as competition regulations necessitate evidence of an anti-competitive motive rather than mere behavioral parallelism. In certain instances, bilateral transparency or unilateral declarations might be perceived as evidence of facilitating practices, potentially impeding effective market competition.

Artificial transparency can be interpreted through industrial economics and game theory prisms. In narrowly oligopolistic markets, where a tacit collusion equilibrium might spontaneously materialize, curtailing uncertainty regarding market player future strategies can markedly truncate the trial-and-error phase and fortify the equilibrium by facilitating the identification of the coordination focal point (OECD 2012).

The ramifications of transparency, be it stemming from information sharing or voluntary algorithmic process disclosure, can be analyzed in the "cheap talk" context. Cheap talk pertains to non-binding signals or communication that can sway others' behavior. The veracity of information provided by the leader can be elusive to ascertain, and it's symmetrically in competitors' interest to adopt a non-cooperative strategy based on the leader's commitment if the leader is perceived as credible. The capacity of algorithms to analyze vast market data and discern conformity patterns to a potential collusion equilibrium can bolster the plausibility of a cooperative equilibrium, inclusive of a collusive one.

It's paramount to recognize that evaluating these factors necessitates meticulous assessment of specific circumstances, market conditions, and the underlying intentions of market participant actions. Competition authorities and courts play an instrumental role in analyzing information exchange, transparency, and algorithmic behavior impacts to ascertain if they constitute anti-competitive practices contravening competition law.

A particularly dynamic field of economic literature has been developing in recent years, extending the scope and effectiveness of econometric tools for detecting abnormal market configurations using machine learning tools. This is the case for the detection of bid rigging in public procurement (Walliman et al., 2023), favoritism in public procurement (Gallego et al., 2021; Decarolis and Giorgiantonio, 2022), or also manipulations in electricity wholesale markets (Brown et al., 2023; Sun et al., 2022).

In the context of enhancing market regulation in the era of complex algorithms, it is essential to focus on equipping authorities with effective market supervision tools. This includes providing them with the necessary technical means and imposing ex ante obligations akin to the practices in the financial sector, such as the registration of prices and conservation of orders. The utility of such practices in the financial world, where detailed records of transactions enable thorough ex-post investigations, suggests their potential effectiveness in monitoring algorithm-driven markets.

Furthermore, the imposition of specific obligations on firms regarding the risks associated with algorithms is crucial. Firms should be required to adopt a comprehensive compliance strategy. This strategy would entail not only ex ante certification of their algorithms but also a commitment to ongoing supervision of their effects, possibly through regular audits. This ensures continuous responsibility and awareness of the impacts of their algorithms.

Additionally, contemplating a shift in the burden of proof, adopting a 'comply or explain' logic, is significant. This approach has parallels in the financial sector, particularly in regulatory compliance, where firms must adhere to regulations or explain their deviations. Translating this logic to the realm of algorithm-driven market practices could prompt firms to proactively ensure that their algorithms do not inadvertently facilitate anti-competitive behaviors. Adopting such measures would represent a substantial advancement in market regulation, addressing the unique challenges posed by algorithmic decision-making in the digital age.

IV. Discussion: Navigating the Complexities of Algorithmic Collusion for Antitrust Enforcers with the Aid of Bandit Algorithms

The proposition of whether a unilateral algorithmic signal can be construed as constituting competitor communication is intricate. While there's no direct bilateral competitor communication in the conventional sense, one firm's algorithmic actions can still transmit a signal to other market competitors. Evans, Tucker, et Fels (2020) examines the problems posed by algorithmic collusion and discusses potential countermeasures, including the interpretation of unilateral algorithmic signals as a form of communication.

In the tacit collusion context, an operator's unilateral decision that contravenes its individual profit maximization and lacks economic rationale might be perceived as a collusion-facilitating factor. This premise is anchored in the notion that if the decision aligns with all market participants' collective self-interest if they concurred to act similarly, but is antithetical to their self-interest if acted upon individually, it might indicate a coordinated strategy. Ivaldi and al. (2007) discuss the economics of tacit collusion, including the role of unilateral decisions that may serve as signals to other firms in the market.

However, the signal's efficacy is contingent on whether competitors' algorithms can capture and interpret it. If competitors possess advanced algorithms capable of real-time subtle signal detection and analysis, the unilateral signal might not remain undetected. In such scenarios, the signal can be as impactful as a financial market order containing concealed information.

The deployment of intricate algorithms and data analysis techniques introduces challenges in determining communication extent and competition law implications. The ability of algorithms to capture and interpret signals might necessitate meticulous scrutiny by competition authorities to evaluate their potential anti-competitive effects and ascertain if they align with the communication scope that abets collusion.

While unilateral algorithmic signals don't encompass direct competitor communication, their influence on market behavior and the capacity of competitors to capture and interpret them are factors warranting consideration in collusion analysis and competition law enforcement endeavors. Upon meticulous examination of the extant literature pertaining to algorithmic collusion and integrating insights from the field of computer science, it appears premature to recalibrate antitrust legislation to address the complexities associated with self-learning algorithms that may engage in collusion within actual market environments. Nonetheless, there exists a subset of algorithmic collusion—specifically, hub-and-spoke configurations enabled by centralized pricing algorithms—which may indeed necessitate prompt legislative intervention. Legislative intervention in such cases may be considered necessary to prevent anti-competitive outcomes that can arise from these algorithmic arrangements. The challenge for lawmakers and regulators is to craft legislation and guidelines that can effectively address the nuanced and technical nature of such collusion without stifling innovation and the legitimate use of algorithms to improve market efficiency.

The phenomenon of algorithmic collusion presents intricate challenges for antitrust enforcement, necessitating a nuanced understanding of its multifarious dimensions. The implications of algorithmic signals and the potential perils of anti-competitive conduct warrant meticulous consideration.

One of the pivotal tools in this exploration has been the utilization of bandit algorithms. Bandit algorithms, traditionally employed in decision-making scenarios where there's a trade-off between exploration (trying out new strategies) and exploitation (sticking with known strategies), offer a pertinent model for understanding the dynamics of algorithmic collusion. Their ability to adaptively learn and make decisions based on partial feedback makes them a representative model for how firms might adjust their pricing strategies in response to the actions of competitors. Directly, they provide insights into how algorithms can iteratively adjust to market signals, and indirectly, they shed light on the broader challenges of decision-making in uncertain environments.

To begin with, competition regulators must cultivate robust methodologies to discern and scrutinize price signals emanating from algorithms. Drawing parallels with financial market oversight, it becomes imperative for authorities to pinpoint subtle signals and anomalous market price patterns. This mandates relentless surveillance and the capability to identify departures from anticipated market dynamics.

Subsequently, the onus of justification might be transitioned to enterprises employing algorithmic pricing paradigms. In scenarios where market conduct diverges from a predetermined model or manifests atypical price configurations, enterprises might be compelled to elucidate their maneuvers. Through rigorous data analysis from these firms, competition regulators can identify anomalies, subsequently placing the onus on enterprises to validate that their algorithmic operations don't culminate in anti-competitive consequences. Nevertheless, a mere inversion of the burden of proof might not be adequate to ascertain anti-competitive conduct, as the motivations and repercussions of algorithmic determinations warrant thorough assessment.

Furthermore, delineating between logical parallel pricing conduct and collusive configurations becomes increasingly intricate, especially within the ambit of unilateral communications. Unilateral signals, irrespective of their digital or traditional nature, engender complexities for both competition enforcers and corporate entities. Firms procuring insights regarding prospective competitor prices might grapple with distancing themselves from the tacit collusion proposition. This complexity is accentuated in the realm of AI-driven automated online pricing technologies, where pricing determinations are predicated on perpetual competitor price surveillance and automated recalibrations. The lucidity of pricing algorithms oscillates between being a proponent of competition and a potential collusion enabler, necessitating a judicious equilibrium.

Moreover, the facet of price transparency emerges as pivotal. Digital pricing strategies, frequently targeting price differentiation, are occasionally perceived as deliberately obfuscated. While competition regulators might advocate for transparency, unveiling pricing algorithms could inadvertently abet collusion. Striking an optimal balance between transparency and collusion deterrence is of paramount importance.

In addressing the intricacies of algorithmic collusion, preemptively thwarting it without compromising algorithmic efficacy becomes a daunting endeavor. Instituting proactive regulations that curtail information assimilation and signal generation capabilities might inadvertently impede pricing algorithms, thereby stifling innovation. Hence, the emphasis should be on specific attributes of information dissemination and assimilation that might bolster collusive equilibria, all while safeguarding the merits of algorithmic precision.

To conclude, the ascendancy of algorithmic collusion introduces profound challenges for antitrust enforcers. The tasks of detecting and interpreting algorithmic signals, discerning anti-competitive intentions and ramifications, and calibrating an equilibrium between transparency and collusion deterrence emerge as pivotal considerations in navigating these challenges. The insights derived from bandit algorithms further underscore the complexities and potential pathways in understanding and addressing algorithmic collusion.

References

1. Abrantes-Metz, R. M., Kraten, M., Metz, A. D., Seow, G. S.. 2012. Libor manipulation? *Journal of Banking & Finance*. Volume 36, Issue 1. https://doi.org/10.1016/j.jbankfin.2011.06.014.
2. Agrawal, S. and Goyal, N.. (2012). Analysis of Thompson Sampling for the Multi-armed Bandit Problem. *Proceedings of the 25th Annual Conference on Learning Theory*, in *Proceedings of Machine Learning Research* 23:39.1-39.26 Available from https://proceedings.mlr.press/v23/agrawal12.html.
3. Arena, Lise; Oriol, Nathalie; and Veryzhenko, Iryna (2018) "Too Fast, Too Furious? Algorithmic Trading and Financial Instability," Systèmes d'Information et Management: Vol. 23 : Iss. 2 , Article 4. Available at: http://aisel.aisnet.org/sim/vol23/iss2/4
4. Assad, S., Calvano, E., Calzolari, G., Clark, R., Denicolò, V., Ershov, D., ... Wildenbeest, M. (2021). Autonomous algorithmic collusion: economic research and policy implications. *Oxford Review of Economic Policy, 37*(3), 459-478. https://doi.org/10.1093/oxrep/grab011

5. Assad, S., Clark, R., Ershov, D., & Xu, L. (2023). Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market. *Journal of Political Economy*. https://doi.org/10.1086/726906
6. Auer, P., Cesa-Bianchi, N., & Fischer, P. (2002). Finite-Time Analysis of the Multiarmed Bandit Problem. *Machine Learning*, *47*(2), 235-256. https://doi.org/10.1023/A:1013689704352
7. Axelrod, R., Hamilton, W. D., The Evolution of Cooperation. *Science*. 211,1390 1396(1981).DOI:10.1126/science.7466396
8. Babaioff, M., Dughmi, S., Kleinberg, R., & Slivkins, A. (2015). Dynamic Pricing with Limited Supply. *ACM Transactions on Economics and Computation*, *3*(1), 1-26. https://doi.org/10.1145/2559152
9. Biais, B., Hombert, J., & Weill, P.-O. (2014). Equilibrium Pricing and Trading Volume under Preference Uncertainty. *The Review of Economic Studies*, *81*(4), 1401-1437. https://doi.org/10.1093/restud/rdu008
10. Borenstein, S., Kwoka, J., & White, L. (2004). Rapid price communication and coordination: The airline tariff publishing case (1994). In *The Antitrust Revolution: Economics, Competition, and Policy* (Vol. 4). Retrieved from https://global.oup.com/us/companion.websites/fdscontent/uscompanion/us/pdf/kwoka/9780195322972_09.pdf
11. Brown, D. P., Eckert, A., & Silveira, D. (2023). Screening for Collusion in Wholesale Electricity Markets: A Review of the Literature. *Working Papers, juin*. Retrieved from https://ideas.repec.org//p/ris/albaec/2023_007.html
12. Bubeck, S., & Cesa-Bianchi, N. (2012). Regret Analysis of Stochastic and Nonstochastic Multi-Armed Bandit Problems. *Foundations and Trends® in Machine Learning*, *5*(1), 1-122. https://doi.org/10.1561/2200000024
13. Calvano, E., Calzolari, G., Denicolo', V., & Pastorello, S. (2019). Algorithmic Pricing: What Implications for Competition Policy? *Review of Industrial Organization*, *55*(1), 155-171. https://www.jstor.org/stable/48702985
14. 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
15. Christie, W. G., & Schultz, P. H. (1994). Why do NASDAQ Market Makers Avoid Odd-Eighth Quotes? *The Journal of Finance, 49*(5), 1813–1840. https://doi.org/10.2307/2329272
16. Cooper, R., DeJong, D. V., Forsythe, R., & Ross, T. W. (1989). Communication in the Battle of the Sexes Game: Some Experimental Results. *The RAND Journal of Economics, 20*(4), 568-587. https://doi.org/10.2307/2555734
17. Crandall, J. W., Oudah, M., Tennom et al. (2018). Cooperating with machines. *Nature Communications*, *9*, 233. https://doi.org/10.1038/s41467-017-02597-8
18. De Coninck, R. (2016). Information Exchanges and Price Signaling: An Economic Perspective. Retrieved from https://www.crai.com/insights-events/publications/information-exchanges-and-price-signaling-economic-perspective/
19. Decarolis, F., & Giorgiantonio, C. (2022). Corruption red flags in public procurement: new evidence from Italian calls for tenders. *EPJ Data Science, 11*(16), 1-38. https://doi.org/10.1140/epjds/s13688-022-00325-x
20. Dorner, F. E. (2021). Algorithmic collusion: A critical review. *arXiv*. https://doi.org/10.48550/arXiv.2110.04740

21. European Commission. (2011). Guidelines on the applicability of Article 101 of the Treaty on the Functioning of the European Union to horizontal co-operation agreements. Retrieved from https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52023XC0721(01)
22. Evans, D. S., Tucker, C. E., & Fels, A. (2020). The Evolution of Antitrust in the Digital Era: Essays on Competition Policy. *Competition Policy International*. Retrieved from https://www.gurkaynak.av.tr/docs/11288-volume-1-the-evolution-of-antitrust-in-the-digital-era-essays-on-competition-policy.pdf
23. Ezrachi, A., & Stucke, M. E. (2017). Artificial Intelligence & Collusion: When Computers Inhibit Competition. *Illinois Law Review, octobre*. https://doi.org/10.2139/ssrn.2591874
24. Gal, M. S. (2019). Algorithms as Illegal Agreements. *Berkeley Technology Law Journal, 34*(1), 67-118.
25. Gallego, J., Rivero, G., & Martinez, J. (2021). Preventing rather than Punishing: An Early Warning Model of Malfeasance in Public Procurement. *International Journal of Forecasting*, *37*(1), 360-377. https://doi.org/10.1016/j.ijforecast.2020.06.006
26. Hansen, K. T., Misra, K., & Pai, M. M. (2021). Algorithmic Collusion: Supra-competitive Prices via Independent Algorithms. *Marketing Science*, *40*(1), 1-12. https://doi.org/10.1287/mksc.2020.1276
27. 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
28. Hingston, P., & Kendall, G. (2004). Learning versus evolution in iterated prisoner's dilemma. In *Proceedings of the 2004 Congress on Evolutionary Computation* (pp. 364-372 Vol.1). https://doi.org/10.1109/CEC.2004.1330880
29. Ivaldi, M., Jullien, B., Rey, P., Seabright, P., & Tirole, J. (2007). The economics of tacit collusion: implications for merger control. *Contributions to Economic Analysis*, *282*, 217-239. https://doi.org/10.1016/S0573-8555(06)82008-0
30. Li, L., Chu, W., Langford, J., & Schapire, R. E. (2010). A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web* (pp. 661-670). https://doi.org/10.1145/1772690.1772758
31. Marty, F., & Warin, T. (2023). Multi-sided platforms and innovation: A competition law perspective. *Competition & Change*, *27*(1), 184-204. https://doi.org/10.1177/10245294221085639
32. Mehra, S. K. (2016). Antitrust and the Robo-Seller: Competition in the Time of Algorithms. *Minnesota Law Review*, *100*, 1323-1375.
33. OECD. (2012). Unilateral Disclosure of Information with Anticompetitive Effects. Retrieved from https://www.oecd.org/daf/competition/Unilateraldisclosureofinformation2012.pdf
34. OECD. (2023). Algorithmic competition - OECD. Retrieved from https://www.oecd.org/competition/algorithmic-competition.htm
35. Page, W. H. (2009). The Gary Dinners and the Meaning of Concerted Action. *SMU Law Review*, *62*, 597. https://scholar.smu.edu/smulr/vol62/iss2/7
36. Salcedo, B. (2015). Pricing algorithms and tacit collusion. *Manuscript, Pennsylvania State University*. Retrieved from http://brunosalcedo.com/docs/collusion-slides.pdf
37. Samà, Danilo. (2014). Cartel detection and collusion screening: an empirical analysis of the London Metal Exchange. SSRN Electronic Journal. 10.2139/ssrn.2425320.

38. Schlechtinger, M., Kosack, D., Paulheim, H., Fetzer, T., & Krause, F. (2023). The Price of Algorithmic Pricing: Investigating Collusion in a Market Simulation with AI Agents. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems* (pp. 2748-2750). https://www.southampton.ac.uk/~eg/AAMAS2023/pdfs/p2748.pdf
39. Schwalbe, U. (2018). Algorithms, Machine Learning, and Collusion. *Journal of Competition Law & Economics, 14*(4), 568-607. https://doi.org/10.1093/joclec/nhz004
40. Sun, B., Deng, R., Ren, B., Teng, M., Cheng, S., Wang, F. (2022). Identification method of market power abuse of generators based on lasso-logit model in spot market. *Energy*. *238*. https://doi.org/10.1016/j.energy.2021.121634.
41. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. MIT Press. Retrieved from https://books.google.ca/books?hl=en&lr=&id=uWV0DwAAQBAJ&oi=fnd&pg=PR7&dq=Sutton,+R.+S.,+%26+Barto,+A.+G.+(2018)).&ots=miwNq4Y2o5&sig=XqkdqYzRVyRTJMaVS5PeLf0--9w.
42. Wallimann, H., Imhof, D., & Huber, M. (2023). A Machine Learning Approach for Flagging Incomplete Bid-Rigging Cartels. *Computational Economics, 62*, 1169-1720. https://doi.org/10.48550/arXiv.2004.05629
43. Waltman, L., & Kaymak, U. (2008). Q-Learning Agents in a Cournot Oligopoly Model. *Journal of Economic Dynamics and Control*, *32*(10), 3275-3293. https://econpapers.repec.org/article/eeedyncon/v_3a32_3ay_3a2008_3ai_3a10_3ap_3a3275-3293.htm
44. Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, *8*, 279-292. https://doi.org/10.1007/BF00992698
45. Xu, X., Lee, S., & Tan, Y. (2023). Algorithmic Collusion or Competition: the Role of Platforms' Recommender Systems. https://doi.org/10.48550/arXiv.2309.14548