3138

REPRINT

Axel Gautier, Ashwin Ittoo , Pieter Van Cleynenbreugel

AI algorithms, price discrimination and collusion: a technological, economic and legal perspective

European Journal of Law and Economics

CORE CENTER FOR OPERATIONS RESEARCH AND ECONOMETRICS

European Journal of Law and Economics
https://doi.org/10.1007/s10657-020-09662-6

AI algorithms, price discrimination and collusion: a technological, economic and legal perspective

Axel Gautier[1] · **Ashwin Ittoo**[1] · **Pieter Van Cleynenbreugel**[2]

Abstract
In recent years, important concerns have been raised about the increasing capabilities of pricing algorithms to make use of artificial intelligence (AI) technologies. Two issues have gained particular attention: algorithmic price discrimination (PD) and algorithmic tacit collusion (TC). Although the risks and opportunities of both practices have been explored extensively in the literature, neither has yet been observed in the actual practice. As a result, there remains much confusion as to the ability of algorithms to engage in potentially harmful behavior with respect to price discrimination and collusion. In this article, we embed the economic and legal literature on these topics in a technological grounding to provide a more objective account of the capabilities of current AI technologies to engage in price discrimination and collusion. We argue that attention to these current technological capabilities should more directly inform on-going discussions on the urgency to reform legal rules or enforcement practices governing algorithmic PD and TC.

Keywords AI · Tacit collusion · Price discrimination · Economics · Competition · Markets · GDPR · Machine learning · Deep learning · Reinforcement learning

JEL Classification C · D · K2 · K4 · L

We thank the handling editor and the reviewers for their helpful remarks and suggestions.This research was funded through the ARC grant for Concerted Research Actions, financed by the French-speaking Community of Belgium.

✉ Ashwin Ittoo
ashwin.ittoo@uliege.be

Axel Gautier
agautier@uliege.be

Pieter Van Cleynenbreugel
pieter.vancleynenbreugel@uliege.be

[1] HEC Liège, LCII, University of Liège, Liège, Belgium

[2] Faculty of Law, LCII, University of Liège, Liège, Belgium

Published online: 14 July 2020

1 Introduction

The advent of artificial intelligence (AI)[1] algorithms has opened up a growing and dynamic area of law and economic research (Van Cleynenbreugel 2020). The increased self-learning and extensive data processing capabilities of these technologies and the potential harm that they can cause to consumers and societal welfare are used as starting points for proposals to modify current legal rules and enforcement practices. It cannot be denied that the increasing amounts of available data coupled with the improvements in AI algorithms can potentially change firms' behavior in the market. Better knowledge about consumers and the use of AI-enabled pricing algorithms may facilitate personalized pricing by firms. In addition, the better observability and predictability of competitors' behavior and the ability of AI algorithms to react almost instantaneously may facilitate coordination on the market. The two phenomena described above are commonly referred to as algorithmic price discrimination (PD) and algorithmic tacit collusion (TC). In recent years, they have not only received significant attention from the research community but have also been in the limelight of mainstream media.

Price discrimination consists in charging different consumers different prices for the same or similar products. Heightened concerns over algorithmic price discrimination practices are related to the massive volumes of personal data disclosed by consumers online. Compared to traditional demographic data, these data could offer a significantly richer trove of valuable information used by sellers to set prices. AI algorithms could exploit these data to generate accurate profiles of consumers and to acquire a deeper understanding of their purchase behaviors (preferences, needs, or dislikes) (Woodcock 2019). Such information is then incorporated in marketing or pricing applications for personalized prices (and personalized product recommendations). It is therefore often posited that the use of AI-enabled pricing algorithms and access to rich datasets of consumer behavior will enable a finer-grained price discrimination. In fact, several models proposed in scientific literature have indeed been shown to be able to price discriminate at a finer degree of granularity (Shiller 2014; Ban and Keskin 2017; Dubé and Misra 2018).

Collusion takes place "when firms use strategies that embody a reward–punishment scheme which rewards a firm for abiding by the supracompetitive outcome and punishes it for departing from it" (Harrington 2018). It has been submitted that AI technologies could have the ability to analyze and oversee the market at an unprecedented scale, scrutinizing competitors' behavior and making rapid price changes to adapt to the environment (Ezrachi and Stucke 2016). In this context, a commonly proposed hypothesis in the competition policy literature is that over a reasonably long time horizon, by repeatedly interacting with each other and their environment, pricing algorithms can learn how to implement collusive strategies based on

[1] Our usage of the term "AI" focuses on machine learning algorithms, which is by far the predominant form of AI deployed in real-life applications. It also includes econometrics and optimization methods, which are closely related to machine learning and have the similar goal of learning from data or their environment.

reward and punishment schemes, leading to supracompetitive prices. Furthermore, these algorithms would be able to sustain collusive outcomes without human pricing intervention. This type of algorithmic collusion, unlike that implemented by traditional cartels, involves no explicit collusive agreement. The algorithms themselves are neither programmed with the intention of collusion nor biased to favor the formation of cartels. Instead, they adopt a collusive strategy through autonomous decision-making. This is what is referred to as algorithmic tacit collusion (Ezrachi and Stucke 2016). Some of the algorithms proposed in the recent scientific literature can effectively learn how to implement collusive strategies over a sufficiently long period (Calvano et al. 2018b).

Even though both PD and TC require the intensive use of data and advanced AI algorithms, they will not necessarily emerge on the same type of market. The economic literature has recognized indeed that PD is not a factor facilitating collusion (Liu and Serves 2007; Colombo 2010; Helfrich and Herweg 2016).

With PD, prices will be closer to the consumer's willingness to pay (WTP). At the intensive margin, consumers will pay a higher price and have a lower surplus. At the extensive margin, there could be demand expansion and the global impact of PD depends on the relative importance of these two effects. With TC, prices will be supracompetitive, increasing firms' profits and reducing consumer surplus. As both practices may affect consumers, they call for closer attention from antitrust enforcers or regulatory bodies.

Although the risks and opportunities of both practices have been explored in the literature and in experiments, neither has been observed in practice (real life). As noted by Schwalbe (2019), most studies have addressed the issues of algorithmic PD and TC predominantly from a legal and economic perspective, treating the algorithms as "mysterious blackboxes". The present contribution ambitions to reach beyond this 'algorithms as mysterious blackboxes' status quo. To do so, we analyze the legal and economic literature on AI PD and TC directly against the background of the evolving computer science literature in the AI-subfields of machine learning and reinforcement learning. This analysis is expected to yield a more objective account of the current capabilities of AI technologies to implement PD and TC, and contribute to better evaluating the urgency of legal or regulatory reform advocated by others.

More particularly, we review the state-of-the-art AI methods that are key to make algorithmic PD and TC a reality. In terms of PD, it cannot be denied that algorithms have been shown to collude and to achieve fine-grained PD in confined, theoretical (research/experimental) settings. However, we also submit that there is a technical chasm between the models proposed in the literature, which can implement finer-grained PD, and the tools currently deployed in practice (Sect. 2). In the same way, the extent to which algorithms, proposed in literature, can collude in real life, and autonomously, largely remains unanswered (Sect. 3).

On the basis of these technological observations, we identify a gap between the theory (models proposed in research) and real life practice of algorithmic PD and TC (Sect. 4). We submit that the presence of this gap needs to be taken into account more explicitly when proposing legal or regulatory reforms. We believe that changes to legal or regulatory frameworks should address the current technological

possibilities rather than being merely speculative. In the context of algorithmic PD and TC, our contribution somewhat counter-intuitively shows that the legal frameworks in place prove sufficiently protective in the current state of technology. In that respect, we call for regulatory reform that is more clearly in line with what the technology one wants to regulate is actually capable of.

2 Algorithmic price discrimination

A fundamental issue that has often been overlooked in debates surrounding algorithmic PD is whether AI algorithms can currently enable finer-grained PD than before. In this section, we question whether there remain technical impediments hindering the implementation of these pricing strategies in real-markets.

To answer this question, we briefly revisit the main economic classifications of price discrimination (Sect. 2.1). We subsequently describe some well-known reported real-life instances where there is strong evidence suggesting that major e-commerce vendors have engaged in PD (Sect. 2.2). To complement these studies, we provide an overview of several recent investigations commissioned by regulatory agencies that experimentally assess the prevalence of PD among e-commerce vendors (Sect. 2.3), and of the mechanisms traditionally used to establish PD (Sect. 2.4).

Next, we review how AI methods have contributed to fine-tuning PD (Sect. 2.5). This analysis enables us to claim that important technical limitations to the deployment of PD models for finer-grained PD still remain in practice (Sect. 2.6).

2.1 Price discrimination in theory

In basic economic theory, two preconditions have traditionally been identified as necessary for price discrimination to take place. First, firms should have the ability to set their prices. This means that a firm must have some market power or at least the possibility of charging differentiated prices (Stole 2003). Second, because goods are sold at different prices, discrimination is sustainable only if consumers cannot engage in arbitrage. The consumer buying the good at a lower price must be unable to resell it to another who is offered it at a higher price (Woodcock 2019).

The economic literature distinguishes three categories of PD depending on the information the firm has on consumers (Belleflamme and Peitz 2015). There is a major distinction between third degree PD (and its extreme form, 1st-degree PD) and 2nd-degree PD. In the former, the firm is able to infer information about consumers WTP from observable and verifiable consumer features; in the latter, the firm cannot and, therefore, has to propose a menu of options, which opens the way to personal arbitrage.

First degree PD (or personalized pricing) corresponds to the case where the firm has access to complete information on each consumer, which enables it to infer the consumer's willingness to pay and to propose a personalized price. Even when having but partial information, the firm can still use this to partition consumers according to their observable features on their willingness to pay (e.g. region, age, gender,

income), and charge different prices to different consumer groups. This form of pricing is known as third degree price discrimination (or group pricing). In both personalized and group pricing, the firm uses the observable information on consumers to tailor its price or, more generally, its offer, and extract more surplus. It is generally acknowledged that more information on consumers allows a finer segmentation of consumers.

Second degree PD is based on a different logic. Here, rather than observing an individual's features, the firm offers different packages (combinations of price and quantity/quality) and the consumers self-select their preferred option.

The use of algorithms, based on AI technologies, could facilitate the emergence of finer-grained PD (tending towards first-degree), or at least help firms to optimize the menu of options that they propose to consumers, as well as the pricing of these options. Questions remain, however, as to what extent current technological possibilities allow for such finer-grained price discrimination to take place. The following sections show that theory and practice remain somewhat disconnected in this respect.

2.2 Real life reported instances of price discrimination in the digital economy

In practice, there exist different tools to implement price discrimination online. First, firms can profile consumers and display different prices accordingly. This practice could be facilitated either by asking consumers to identify themselves, by cookies identifying them or by the caching mechanisms available in most web browsers. (Cahn et al. 2016). This information could be exploited subsequently for personalization, through targeted ads, product/service recommendations and personalized pricing. Information on identified regular consumers, ZIP code or OS type can be used to segment consumers. Alternatively, firms can display a flat price online but offer targeted coupons to consumers. Second, firms can display personalized recommendations and present different price and product combinations to different consumers.[2] There is evidence supporting that rank in algorithmic search influences click-through rates (Ghose and Yang 2009). As consumers often click on the highly-ranked links or products (i.e. those displayed at the top of lists or search results) presented to them, personalized recommendations can be an effective tool to price discriminate between them.

One of the best-known reported PD cases in practice is that of Amazon, in which a user noticed that the price of a DVD offered to him dropped from $26.24 to $22.74 after he deleted cookies identifying him as a regular customer. This "DVD case" created much outrage among other customers. Amazon attributed the price difference to random price tests, and eventually refunded all customers who had paid the higher price (Cavallo 2018). In another similar incident, a consumer observed that mahjong tiles were priced at $79.99 in her shopping cart, when a few minutes

[2] In the case of new products, consumers may not be aware of their willingness to pay. Personalized recommendations (based on recommender systems and AI methods) may help consumers realize their preferences. Marketing techniques, such as A/B testing, can also help in these situations.

earlier, they were on offer at $54.99. Subsequently, upon clearing her browsing history, the product was priced at $59.99 (Townley et al. 2017).

A study by Propublica, a US-based non-profit organization, revealed that The Princeton Review used information on ZIP codes to determine the price of its online SAT courses. Since ZIP codes are often a proxy for ethnicity, this resulted in some ethnic groups being charged a higher price than others. For example, in ZIP codes dominated by a specific ethnic group, the course price ranged between $7 200 and $8 400 respectively, while in other ZIP codes where this same group was a minority (and a different ethnic group was the majority), the course was offered for $6 600.[3] Another case of PD based on ZIP codes involved Staples. As described in a Wall Street Journal (WSJ)[4] report, the Staples e-commerce website randomized between showing a high price and a low price to users based on their ZIP codes. For instance, the same stapler could cost either $15.79 or $14.29, depending on the user's ZIP code. Furthermore, consumers living within a 20-mile radius of rival stores (e.g. OfficeMax and Office Depot) were offered discounted prices. This resulted in consumers from rural areas, most of whom would be worse-off, paying higher prices than their better-off urban counterparts.[5]

Another instance of price discrimination has been the personalization of the product ranking displayed to consumers. Travelocity was found to provide different search results for users of Safari web browsers on iOS devices, who were shown different hotels and in different order from those shown to users of Chrome on Android and Safari on OS X. Furthermore, iOS users were shown prices that were lower (by ~ 15% on Travelocity) than for other users (Hannak et al. 2014). Market analysis of Orbitz, another online travel agent, revealed that Mac users were likely to spend $20-$30 more per hotel night than their PC-using counterparts. Furthermore, Orbitz's analysis showed that the probability of Mac users booking a 4- or 5-star hotel was around 40% higher than the corresponding probability for PC users.[6] Consequently, Orbitz began personalizing search results of Mac users by placing more expensive options higher up in the ranked-lists of hotels. After the case was made public, Orbitz terminated the use of the personalization algorithm (Hannak et al. 2014).

In the retail sector, Home Depot was found to price discriminate against mobile users, who were steered towards more expensive products compared to desktop users. Discrimination was also observed across mobile users. Android users were served with 24 results following a search, while iOS users were served 48 search results. Furthermore, the search results provided to these different groups were also different, with close to zero overlap between them (Hannak et al. 2014).

[3] https://tinyurl.com/h4y6cqp.
[4] https://tinyurl.com/llu4t9d.
[5] https://tinyurl.com/yaryjq76.
[6] https://tinyurl.com/k8hkfgv.

2.3 Experimental assessment of price discrimination practice

Beyond these specific examples, the extent to which PD practices are implemented in practice and on a more general scale remains largely unknown. For this reason, some regulators have tried to assess their prevalence empirically by conducting ghost shopping experiments.

The *Austrian Chamber of Labor* (Arbeiterkammer Wien) conducted a large-scale study to examine the prevalence of PD across 33 e-commerce vendors, including Amazon, Lufthansa, AirBerlin, Austrian, Opodo, Booking.com and Heine. Their e-commerce sites were accessed on various devices, e.g. desktops, laptops and notebooks (Windows, Apple), iPads, and smartphones (including those running on different operating systems) in various places in Austria and Germany. The experiment was conducted in March 2017 and revealed substantial price variations across devices along with price differences across time and between the two countries (Competition and Markets Authority Report, 2018).

The *DG Justice* of the European Union commissioned research investigating the extent of PD in EU member states (DG Justice 2018). The research was a mystery shopping exercise in four online markets (airline tickets, hotels, sports shoes, and TVs). Their study revealed the prevalence of personalized ranking of online offers in more than half of the visited e-commerce websites. They, however, found no substantial evidence of price discrimination across EU member states nor in the four markets considered.

The study in Reinartz et al. (2017) (as part of a larger effort of the *German Expert Council for Consumer Affairs, Federal Minister of Justice and Consumer Protection*) found that PD based on user characteristics (e.g. demographics) and operating systems was apparent for packaged tours at the higher end of the price spectrum. Such price differentiation, though, could be empirically demonstrated neither for lower-priced tours nor for other market sectors (consumer electronics, sporting goods and garden furniture). On the basis of these findings, the Expert Council concluded that PD in practice was "verifiable only in isolated cases".

The United Kingdom *Competition and Markets Authority* (CMA) conducted an investigation, aimed at collecting *prima facie* evidence on the existence of PD. Its analysis particularly determined whether prices varied in response to a number of variables, such as operating systems, logged versus anonymous search, direct versus indirect assess to e-commerce sites, geographical location (within the UK) and past purchasing history. Very little evidence suggesting the practice of price discrimination was found during the experiments. However, variations in search results were observed across different users based on the variables listed above (Competition and Markets Authority Report, 2018).

2.4 Mechanisms of price discrimination

Despite the limited evidence on the existence of PD in real-life markets, there is huge potential for PD based on AI algorithms and data. In this section, we briefly

explain the main techniques for extracting information from consumer data. Algorithms play an important role in price discrimination in helping to cluster customers into homogeneous groups and to estimate their willingness to pay. We briefly explain these two concepts below.

Third-degree PD is made possible by increasingly sophisticated customer partitions into segments or clusters. Datasets often record many consumer features and large databases have millions of consumer records. Many consumers, however, will exhibit similarities so that they can be clustered into segments.

The basic principle of clustering algorithms is to compute the similarity between different data points, and then partition the data points based on their similarity. Two data points whose similarity exceeds a user-defined threshold are assigned to the same cluster.

The objective of clustering algorithms is to generate clusters so as to maximize *intra-cluster* similarity and minimize *inter-cluster* similarity. These criteria are satisfied when each cluster contains data points that are maximally similar to each other, and minimally similar to data points in other clusters. After identifying the clusters, a firm can then determine how much to charge each consumer in each of the different segments.

The aim is to determine whether a consumer, i, is willing to pay a given price t_i, for a given product (or service). The seller can modify the price across consumers or consumer groups. By way of illustration (and for mathematical convenience), we consider a specific case of the discrete choice model, in which consumer responses, y_i, are either "1" or "0" (i.e. a dichotomous choice model), respectively indicating whether they are willing to pay t_i or not.

For each consumer, the firm has access to data pertaining to their socio-economic and demographic background, such as age, gender, region, occupation, income level and education level, i.e. the so-called explanatory variables.

If, we assume that the willingness to pay (WTP) of consumer i can be expressed as a linear model (Lopez-Feldman 2012), we can write:

$$WTP_i\left(z_i, u_i\right) = z_i\beta + u_i$$

In this equation, z_i is a vector of independent (explanatory) variables. The importance of each variable is captured in the coefficient β and u_i is known as an error term (to account for "noise" in the measurement).

It is expected that consumer i will agree to pay a price of t_i if t_i is less than his/her WTP. Specifically,

$$WTP_i > t_i$$

The corresponding probability can be expressed as

$$\begin{aligned} P(y_i = 1|z_i) &= P(WTP_i > t_i) \\ &= P\left(z_i\beta + \mu_i > t_i\right) \\ &= P(u_i > t_i - z_i\beta) \end{aligned}$$

This estimated probability is usually solved using a *probit* or *logit* regression.[7]

To price discriminate, a firm will first generate clusters of similar consumers. As explained above, the similarity is computed on the basis of the distance between the consumers' attributes, encoded in vectors of independent variables, such as age, gender, and income. Once the consumers have been segmented into clusters, the firm then computes the probability that each segment will purchase a product at a given price. Since each cluster consists of consumers with similar attributes (independent variables), their purchasing probability (at a given price) will be very similar or vary within a small range. Conversely, the purchasing probabilities of consumers across distinct clusters will differ. The firm then can charge different prices to the distinct consumer groups, resulting in group pricing. More data and more powerful algorithms obviously lead to finer grained consumer segmentation. In particular, in addition to socio-demographic information, the advent of the web and social media platforms has enabled firms to harvest a trove of customer behavioral information. This can easily be garnered from the customers' online interactions, and be mined to extract pertinent behavioral information, such as their browsing patterns, likes, dislikes and product preferences. The incorporation of behavioral information in more advanced models proposed in the scientific literature yields a more refined estimate of the WTP of consumers.

2.5 Algorithms for price discrimination: overview of the scientific literature

Several scientific studies have investigated the ability of algorithms to offer refined techniques to compute the best prices to apply to consumers or consumer groups. These studies attempt to estimate a more "precise" value for the WTP of more narrowly-defined consumer clusters. To compute such estimates, they exploit not only classical demographic variables but also behavioral ones, such as consumers' browsing behavior, purchasing patterns and preferences garnered from consumers' online activity. This information is fed to machine learning and econometric methods so as to obtain a more refined estimate of purchasing probabilities.

Shiller (2014) proposed a model estimating the probability that an individual subscribes to a specific package on Netflix. The subscription probabilities are inferred using a probit regression from a dataset of nearly 5000 variables, encompassing demographic and behavioral variables. Experimental evaluations revealed that personalizing prices based on demographics alone led to relatively insignificant increases in profits, amounting to 0.8%. Conversely, the jump in profits was more impressive when additional behavioral variables, such as

[7] If instead of a 1 (yes)/0(no) response, consumers can also rank their preferences at a given bid price t_i, then an ordered probit (tobit) can be used. This model can also be extended to rank product alternatives.

frequency of web site visits and time spent browsing, were included in the models. Profits in this case were estimated to be around 12% higher than what could be achieved without PD, with some consumers paying almost twice the price paid by others.

Chen et al. (2015) developed a general model implementing PD. A decision maker observes a set of consumer attributes, known as 'the context'. Based on this context, the decision maker chooses one type of action from a predefined set of possible actions, for instance altering the prices. Once the action is taken, she observes the outcome, i.e. whether a product is purchased or not. Repeating these actions, the objective is to train the algorithm to select the action that maximizes rewards given the context. The interactions between the context, actions and outcomes are learnt by a logit regression, and are expressed as conditional probabilities, i.e. the probability of an outcome given the action chosen and the context. An expected revenue (the reward) is computed from the estimated probabilities. This enables the seller to infer the price that consumers are willing to pay for products, to charge these consumers accordingly, and subsequently to maximize the rewards. Experimental evaluations performed on both simulated data and real-life transactional data (airline tickets) showed that the proposed model for determining the price consumers are willing to pay, and charging them accordingly, led to an increase in revenue compared to a single best price strategy.

The model for personalized pricing presented in Ban and Keskin (2017) operates on the basis of the same premises. The proposed model learnt the impact of consumer features on product demand and used this information in pricing decisions to maximize revenue over time. The relationship between product demand and pricing was established using a lasso regression method. The authors showed that their approach to personalized pricing achieved the best performance compared to other standard classical pricing strategies, such as myopic pricing, pricing for the average customer, and segment-then-optimize policies akin to third-degree PD based on consumer clusters.

Dubé and Misra (2018) posited that consumers' price sensitivities could be characterized by a vector of observable attributes (features). The goal for a seller was then to statistically learn about demand from heterogeneous consumers. The authors proposed a method for an extreme (refined) form of third degree PD. Two machine learning methods were employed in this study: a lasso regression to select only those consumer features that have a statistically significant influence on demand, and a logistic regression for approximating the firm's uncertainty about demand. The proposed model was implemented and evaluated in a business-to-business pricing application that optimizes the price charged to each prospective new consumer of a recruiting firm. The application was trained using experimental data to learn the relationship between price sensitivity and consumer characteristics. The trained model was then applied to a sample of new consumers, and the benefits of using the recommended prices (from the model) were evaluated against a status-quo price. The model showed that personalized prices could lead to large profit increases.

2.6 Is AI enabling finer-grained PD?

From the foregoing discussion, there is no strong evidence showing that firms are actually implementing finer-grained PD in practice. From the reported PD cases and market investigations conducted by various agencies (Sects. 2.2 and 2.3), it appears that only traditional forms of PD, based on socio-demographical characteristics (e.g. OS/device, region), are common. From a technological perspective, such PD implementations do not require sophisticated AI-based methods. Instead, they can be straightforwardly implemented using traditional clustering algorithms or rule-based (IF–THEN-ELSE) methods (e.g. IF OS = Mac THEN price = price * 1.5).

Conversely, experiments reported in scientific studies (Sect. 2.5) have shown that the use of sophisticated machine learning/econometric methods coupled with consumer data, and in particular, data pertaining to behavioral information, can enable firms to price individual consumers more accurately. This results in a finer-grained PD and is often translated into growth in expected profits. Consequently, there still remains a gap between PD implementations in practice and the models developed in scientific experiments.

3 Algorithmic tacit collusion (TC)

Another phenomenon that attracted the attention of policy makers and scholars is that of algorithmic tacit collusion (TC). In this section, we provide an overview of the recent literature based on state-of-the-art machine learning methods, and offer a scientific account of the propensity of algorithmic agents to form collusive agreements. Our analysis leads us to claim that there remain important limitations hindering the deployment, in practice, of advanced machine learning methods for algorithms to learn how to collude autonomously.

TC takes place when AI algorithms learn how to collude in an autonomous or near-autonomous manner, involving minimal or no human intervention (Sect. 3.1). A conjecture put forward in several pieces of legal scholarship in this respect is that advances in (machine learning) algorithms and their prevalence in pricing applications, coupled with the availability of massive volumes of data, will render TC sustainable in markets even without significant levels of oligopoly concentration (Ezrachi and Stucke 2017) (Sect. 3.2). However, to date, while some scientific studies have shown that algorithms can collude or facilitate collusion under specific constraints, there have been no real-life reported cases of algorithmic tacit collusion (Sect. 3.3). As a result, there exists a gap between theory and practice in the context of TC as well (Sect. 3.4).

3.1 The economics of tacit collusion

In a static environment, the market equilibrium is a Nash equilibrium where each firm's strategic choice (price, quantity, etc.) is the best reply to the competitors' own.

The competitiveness of the market equilibrium depends on specific factors, such as the nature of competition (price vs. quantity), the number of firms, and the degree of product differentiation. The idea behind any kind of collusion (explicit or tacit) is that a collective and coordinated deviation from the Nash behavior can lead to higher profits. If such a deviation from the static Nash behavior is achieved by an explicit coordination by firms, then colluding firms are part of a cartel, which is prohibited under competition law.

Alternatively, coordination can be achieved without communication, and it can emerge as an outcome of a repeated game. In this case, collusion is tacit. To be sustainable, tacit collusion requires some ways of explicit coordination, for instance, information exchange, to monitor the competitors' behavior.

Collusive equilibria in repeated games are supported by grim-trigger strategies of the like "player i plays the collusive outcome at period t, if all the other players have played the collusive outcome at all the previous periods and returns to the static Nash behavior otherwise". This equilibrium is based on threats and rewards. If a firm adopts "good" behavior by playing the collusive outcome, it will be rewarded as collusion will continue for the next period (the carrot); if a firm adopts a "bad" behavior by deviating (profitably) from the collusive outcome, it will be punished as the industry will return to the static Nash equilibrium (the stick). A credible punishment mechanism is the cornerstone of a collusive outcome.[8] It must be sufficiently severe to offset the benefits of a unilateral deviation and it must be optimal for the firms to carry on the punishment when a firm deviated (Garces-Tholon et al. 2009).

Supported by such strategies, tacit collusion can emerge as a subgame perfect Nash equilibrium in an infinitely repeated oligopoly game. The literature on tacit collusion has focused on the conditions guaranteeing that a collusive equilibrium exists and is stable (d'Aspremont and Gabszewicz 1986), while the policy literature has focused on the facilitating factors, such as symmetry between firms, market transparency, multi market competition, and likelihood of entry (Motta 2004, chapter 4).

Empirically identifying tacit collusion is a challenging task as, contrary to cartels, there is no explicit legal evidence to attest its existence. Furthermore, it is difficult to infer collusion from market data. For this reason, there is little evidence documenting the importance of the phenomenon.

The experimental economic literature shows that collusive outcomes are hard to achieve even in laboratory experiments. In particular, Huck et al. (2004) performed a meta-analysis covering 20 laboratory experiments of Cournot games. They showed that coordination is rare in markets where there are more than two firms. Fonseca and Normann (2012) reported similar results for Bertrand oligopolies with prices close to marginal cost for four firms, despite a general belief that collusive outcomes are more likely in Bertrand type of games than in Cournot types (Suetens and Potter 2007). Cost asymmetries make collusion even less likely (see for instance Fischer and Normann 2019). Many experiments have shown that firms fail to coordinate

[8] A similar mechanism of threats and rewards is also necessary to implement an explicit coordinated outcome as cartels are fundamentally unstable.

without pre-play communication (Fonseca and Normann 2012). The empirical and experimental literature corroborates the idea that coordination without communication is far from being obvious, even in laboratory experiments where many environmental variables are controlled for.

3.2 Algorithmic collusion

The term "algorithmic collusion" lends itself to several interpretations, and a clear distinction needs to be made between two situations. The first situation concerns *algorithmic explicit collusion,* where algorithms implement an existing collusive strategy potentially defined or agreed upon by humans. A classic example of algorithmic explicit collusion is the "Amazon poster case".[9] In this case, Trod Ltd had agreed with GB Eye not to undercut prices for posters and frames sold by each party on the Amazon Marketplace (via Amazon UK) from 24 March 2011 to 1 July 2015. This agreement was implemented using automated pricing software. Trod and GB Eye also communicated regularly to ensure that the agreement was stable and to deal with technical issues pertaining to the pricing software. To do so, the algorithms had been programed to react in a predictable manner with the view of enforcing a cartel agreement, conceived by humans.

The second situation concerns *algorithmic tacit collusion*. In this situation, algorithms did not need programing explicitly to enter in collusion. Nevertheless, collusive arrangements may emerge tacitly without any form of explicit agreement from the algorithms' actions, such as by adjusting prices in trying to optimize a profit-making objective. In this situation, human intervention is minimal or completely inexistent. Ezrachi and Stucke (2015) elaborate on three schemes via which algorithms could tacitly collude: hub and spoke, predictable agent and autonomous agent.

In the hub and spoke model, several competing firms (the spokes) delegate their pricing decision to the same intermediary (the hub), who can coordinate the prices of the spokes. We do not consider such price coordination as *algorithmic tacit collusion* for the following two reasons. First, the spokes would have to collude by coordinating on their choice of the hub. What could happen is that even if there is no explicit coordination within the hub, the hub and spoke structure could lead to price parallelism, as the hub would use the same market data and algorithms for the pricing decisions of all its client firms. However, this would likely require some kind of communication transforming the hub and spoke scenario into one of explicit collusion. No new competition law problems would emerge (Schwalbe 2019). Second, if there is a single large pricing hub, explicit price coordination by the hub would be considered an abusive practice falling under Article 102 TFEU or as a decision taken by an association of undertakings falling within Article 101 TFEU. In both cases, the algorithm plays a minor role. As stated in Ohlhausen (2017), "If the word

[9] https://tinyurl.com/y8ry62b9 and https://tinyurl.com/y93ueqvf.

'algorithm' can be replaced by the phrase 'a guy named Bob'", then algorithms can be dealt with in the same way as in a traditional hub-spoke cartel.

The predictable agent scheme involves the use of pricing algorithms, which react to external factors, but in a predictable manner. These algorithms have limited ability to learn and to act autonomously since their responses are often based on hard-coded rules, as in "win-continue lose-reverse" or "match low price/tit for tat" algorithms. An algorithm implementing the predictable agent scheme is that proposed in Zhou et al. (2018), which relies on a zero-determinant (ZD) strategy in an iterated prisoner's dilemma (IPD) setting. ZD strategies are a type of probabilistic and conditional strategies, which makes it possible to unilaterally set the expected payoff of an opponent in an IPD setting, irrespective of the opponent's strategy (coercive strategies), or else to set the ratio between the player's and their opponent's expected payoff (extortionate strategies) (Adami and Hintze 2013). The outcomes of such coordination remain unclear (Schwalbe 2019). There is a potential risk of the algorithms spiraling out of control, as illustrated by the "Making of a Fly" textbook case, in which constant adjustments by the pricing algorithms of competitors caused the book's price to skyrocket to more than $23 million.[10]

The autonomous agent scheme involves pricing algorithms, designed to maximize (or minimize) objective, such as profits (or losses). If these algorithms are based on sophisticated machine learning paradigms, then they can experiment with various actions (e.g. increase prices) to determine the optimal policy, which satisfies their objectives. Collusion emerges when the algorithms determine that their (joint) profit is maximized when they coordinate their prices in the absence of any predefined agreement and this outcome is supported by a credible punishment mechanism.[11] This scenario exhibits a "higher degree of fit" with the notion of algorithmic tacit collusion (TC). It also implements "complete" TC, alleviating the need for human intervention or the injection of human knowledge in its operations. A key question to be asked, however, is the extent to which such a scenario is likely or even possible from a technical point of view.

3.3 The likelihood of algorithmic tacit collusion in the autonomous agent scenario

Against the theoretical background sketched above, this section reviews recent studies investigating whether algorithms can learn how to collude tacitly, and aims at a deeper understanding of the underlying mechanisms of algorithmic tacit collusion. We do not include recent research on the emergence of collaboration and competition between agents (e.g. Leibo et al. 2017; Lerer and Peysakhovich 2017;

[10] https://tinyurl.com/yd2ygk84.

[11] Supracompetitive prices are not enough to attest the existence of TC. Brown and MacKay (2019) show that the use of algorithms for pricing modifies the nature of the pricing game. Through the possibility of reacting almost instantaneously to changes in the environment, the use of algorithms by some firms changes a simultaneous pricing game into a sequential one, leading to higher prices. There, however, is no collusion behind that.

Peysakhovich and Lerer 2017; Crandall et al. 2018) as they do not specifically address TC. Nor do we include studies based on finite state automata, as their popularity has gradually waned over the years, and they have been replaced by and outperformed by more modern approaches, including those based on machine learning. For an overview of finite state automata methods, we refer the reader to Schwalbe (2019). Our focus, instead, is on studies investigating AI techniques for algorithmic collusion. The predominant technique in these studies is reinforcement learning.

Reinforcement learning (RL) exhibits several properties that make it especially convenient to study the interaction of profit maximizing algorithms (Ittoo and Petit 2017). First, it is a based on trial-and-error, involving a pricing algorithm experimenting with its environment and various strategies before finding the optimal one. Second, there is an exploration versus exploitation tradeoff, enabling algorithms to mimic pricing tradeoffs faced by oligopolists. In addition, model-free RL methods alleviate the need for models of the markets, e.g. demand functions or competitors' cost functions. Over the years, RL has become a de facto approach for investigating TC.

Q-learning is a popular (model-free) RL method. It involves an algorithmic agent (agent) in a given environment, consisting of several states. In each state, the agent performs an action according to its strategy (for the current state), and receives a corresponding reward. In the successor state, the agent then chooses the best action. The action is selected according to an action-value function, which is subsequently updated depending on the reward received.

One of the first studies investigating algorithmic collusion is that of Waltman and Kaymak (2008). They showed that Q-learning agents (algorithms) could tacitly collude in the setting of a Cournot oligopoly. Later studies, however, and that of Izquierdo and Izquierdo (2015) in particular, demonstrated that the results achieved by Waltman and Kaymak (2008) were brittle, and broke down in the presence of small perturbations in costs, prices and other parameters. More recently, Klein (2018) demonstrated that simple Q-learning agents could learn how to collude by achieving supracompetitive profits in stylized environment with sequential prices. In their first set of experiments, a Q-learning agent played against a fixed strategy, implemented as a myopically optimizing competitor that always undercut. Initially, both algorithms (the Q-learning agent and the competitor) benefited from the exploration strategy of the Q-learning algorithm. This mutual benefit, however, gradually diminished as the Q-learning algorithm acquired a more informed view of the environment over time, and started to optimize its instantaneous profits. Eventually, both algorithms learnt to coordinate their behavior so as to exploit the dynamics of Edgeworth price cycles, resulting into substantially higher profits than under competition. In the second part of the experiments, two Q-learning algorithms competed against each other. Similarly, they learnt to coordinate their behavior and achieved even higher profits.

In their preliminary experiments, Calvano et al. (2018a) pitted two Q-learning algorithms against each other to determine if they reached a cooperative equilibrium in a stage game, where the unique Nash equilibrium is to "choose low price". Their environment was designed to capture the idea that for a collusion to be sustainable, firms should resist their temptation to undercut their rivals and to achieve short-term

profits. Their results revealed that after several iterations, both Q-learning algorithms achieved a high level of collusion. Reported profits gains were 70% of the maximum possible gains after 1000 iterations. The gains increased further with a large number of iterations, amounting to 85% (of the max. gain) after 100 000 iterations.

In their subsequent study, Calvano et al. (2018b) provided additional evidence that Q-learning algorithms could systematically learn to play sophisticated collusive strategies. They particularly allowed several Q-learning algorithms to interact in a Bertrand oligopoly setting, with logit demand and fixed marginal cost. Despite the absence of theoretical guarantees that convergence can occur in the presence of more than two agents, their experiments revealed that convergence was achieved in 99.9% cases. Furthermore, the algorithms charged supracompetitive prices, resulting in supracompetitive profits. Prices were not as high as they would have been in a monopoly, but were superior to those under the Bertrand-Nash equilibrium. In addition, price dispersion was low and firms tended to price symmetrically. They also simulated cheating by manually lowering the price of an algorithm over a given time period. Interestingly, a punishment mechanism, as in classical collusive agreements emerged, i.e. algorithms punished defections. The punishments inflicted were proportional to the extent of the deviations and lasted for a finite time period, before reverting to prices close to the pre-deviation prices.

3.4 Is AI enabling TC?

There is scant evidence that, in practice, firms are implementing pricing algorithms that can learn autonomously how to collude. Among the three forms of algorithmic collusion presented by Ezrachi and Stucke (2015), we discussed how only the autonomous agent actually qualified as an instance of TC. Instead, the hub and spoke and predictable agent forms qualify as examples of explicit collusion. The main (if not the only) case of algorithmic collusion reported to date, the "Amazon poster case", is one of explicit (and not tacit) collusion as the algorithms were programed to execute a set of predictable actions (e.g. not to undercut competitors' prices). At the same time, recent scientific results have shown that algorithms can learn how to collude and implement punishment mechanisms in an autonomous manner. Despite these theoretical possibilities, their implementation in practice still lags behind what turned out to be possible in an experimental setting.

4 Technical challenges and their impact on economic and legal policymaking

In the previous sections, we described how algorithms could implement (i) finer-grained PD and (ii) TC. Yet, at the same time, we also showed that these forms of PD and TC are not (yet) likely to have been implemented in practice. This could be attributed to the gap between experimental research and experimental settings in which these algorithms have been developed and evaluated, and the real-world

environments where they are to be implemented in practice. In most cases, the sophisticated algorithms developed in the scientific literature cannot be directly transposed to real-world environments due to various challenges and constraints present in such environments.

In this section, our aim is to analyze and explain the gap between theory and practice by focusing on current technological challenges and limits in terms of machine learning. These limits, we submit, need to be embedded more firmly in current economic and legal policy debates. Such embedding calls into question the urgency of legal reform. Moreover, the regulatory frameworks in place are seemingly quite ready to accommodate PD and TC methods that can be implemented given the current state of technological developments.

We focus on the technical challenges causing or maintaining the gap between theory and practice for algorithmic PD (Sect. 4.1) and TC (Sect. 4.2). We also discuss the impact of these challenges on economic and legal policies currently in place or under consideration. In light of these observations, we are able to offer an overview of the differences between and the challenges encountered by both algorithmic PD and TC (Sect. 4.3) before offering our general conclusions in this respect (Sect. 4.4).

4.1 Price discrimination: technical challenges and economic and legal policy responses

Despite the real possibilities for price discrimination, there remain various technical challenges (Sect. 4.1.1). These, together with the economic and legal frameworks currently in place, are likely to limit the real-life emergence of algorithmic price discrimination (Sect. 4.1.2).

4.1.1 Technical challenges

The models for PD proposed in the scientific literature do not adequately deal with a number of longstanding algorithmic challenges. The first one pertains to the high-dimensionality of the data. Consumer features in the proposed PD models are encoded in high-dimensional vectors, with each dimension (vector element) corresponding to an attribute (e.g. age, price paid in the past). These vectors tend to be sparse, with a huge proportion of elements having "0" values (as it is not always possible to collect information on all attributes on all consumers). However, processing such high-dimensional vectors is computationally expensive, requiring huge amounts of memory and computational time—an issue referred to as the "curse of dimensionality".

Furthermore, the sparse nature of the vectors makes it difficult to accurately assess the impact of specific consumer features on demand. Ban and Keskin (2017) completely overlooked the issue of high-dimensionality. Shiller (2014) and Chen et al. (2015) addressed it using a stepwise feature selection procedure, involving the gradual incorporation of additional features in the models until a drop in performance (predictive accuracy) was observed. These methods, though, tend to be computationally expensive. It is doubtful whether such procedures will scale up in real-life, dynamic scenarios. Dubé and Misra (2018) addressed this issue using the lasso

regression to select the most predictive features. Lasso performs feature selection by pushing many of the coefficients (associated to specific features) to zero. For some applications, however, this approach of reducing coefficients to zero might prove too drastic and could discard otherwise meaningful features. Furthermore, given a set of highly correlated features, lasso tends to select only one of the features, discarding all others. An alternative to the lasso could be the elastic net regression, which performs less drastic penalization and does not suffer from the issue of correlated features.

Another algorithmic shortcoming pertains to the supervised learning paradigm employed by theoretical PD models. Supervised learning involves training a machine learning algorithm on past data, which have been annotated/labeled with the information of interest, such as whether a customer has purchased a product at a given price. However, with the possible exception of pre-processed datasets purchased from brokers, most data, generated by or harvested from common online sources are originally unlabeled. Subsequently, they have to be manually prepared and annotated to make them amenable for supervised learning. As can be expected, such a manual annotation is costly and error-prone.

One approach that could be adopted to alleviate the aforementioned issue in acquiring information on consumer WTP is auction mechanisms, in which bidders (consumers) explicitly reveal the price they would be willing to pay for a product. This information, together with other available consumer features, would constitute a reliable source of labeled data for training supervised learning methods. Varian (2009) adopted this approach to determine the value that advertisers are willing to pay in the context of search engine advertisements (SEA). However, even though auction mechanisms represent an efficient means to acquire labeled data, the data would be limited in scope as it would be applicable only to those consumers that have participated in online auctions (with eBay as the prototypical example). Furthermore, to maintain its competitive advantage, an auctioneer (auctioning platform) might not disclose or even sell information acquired about its consumers. Therefore, firms would have to implement their own auctioning mechanisms in order to obtain information on consumers' WTP (or pay a high price to acquire such information). Also, the risk with auctions is that consumers generally bid at prices below their WTP. Machine learning models trained on such data would fail to predict the true WTP of consumers, which is more relevant for firms as it ensures higher profits.

An alternative solution to supervised learning could be reinforcement learning (RL). RL presents an attractive basis for pricing models as it does not require any explicit knowledge of consumer purchasing behavior. Rana and Oliveira (2015) proposed an RL approach for *dynamic pricing* at peak/off-peak hours or at specific days of the week for interdependent products. Their RL approach was implemented using Q-learning, where states correspond to the available inventory for each product (service), actions correspond to price changes and state transitions specify how the actions affect transitions from one state to another. Experimental evaluations revealed that the RL-based approach for dynamic pricing increased revenue compared to a fixed pricing policy. Furthermore, the proposed model learnt good policies (revenues started to increase) after a reasonable number of (training) episodes. The authors did not specify the actual time taken to arrive at these policies, and whether the training time would be acceptable in practice.

An overlooked aspect in the proposed models concerns their scalability. Most models are based on machine learning techniques, which at their core are optimization problems (e.g. finding the set of consumer features to optimally price discriminate). Mathematically, these problems cannot be solved in polynomial time (known as NP-hard problems approximated using heuristics). Nevertheless, determining the best heuristics is challenging as it involves searching the space of all possible combinations of (consumer) features (and corresponding values), and selecting that combination which maximizes some objective (or minimizes an error). As can be expected, this issue is compounded in real world datasets. These datasets are larger, noisier and contain many more features than the experimental ones on which the models were evaluated (Lee et al. 2006). Furthermore, current studies report neither on the models' run time, and in particular, on the time taken to train (i.e. to fit) the models over the available (experimental) datasets, nor on the necessary pre-processing steps required to transform available data into a format amenable for their (computational) analysis.

One solution to avoid these issues is to rely on cloud platforms offering "machine learning pipelines as a service", such as Amazon Web Services,[12] Microsoft Azure,[13] Google[14] or SalesForce CRM.[15] While relying on cloud computing services offers several advantages, it also entails a number of drawbacks. Firms would have to delegate sensitive information, such as pricing strategy and consumer data to third parties. Another issue could be the cost of subscribing to a cloud provider and using the available machine learning services. This cost could vary according to the usage frequency, the data volume processed or the amount of processor time consumed. In addition, the solutions offered by most cloud providers are generally limited to classical machine learning applications, such as customer segmentation (which is at the basis of 3rd-degree PD). The advanced algorithms for finer-grained PD, as developed in the scientific literature, will typically not be available as cloud services. It is also worth mentioning that these cloud service providers are often platforms selling or providing a number of additional services (Google, AWS), often competing with other firms on specific markets (Amazon marketplace or Google Flights). It is highly unlikely that these other firms would allow their (much larger) competitors to access their data and to process it using their machine learning services.[16] It is therefore reasonable to expect that while large companies with a solid digital backbone and advanced IT infrastructure (e.g. Google, Amazon) would be able to adequately address the scalability issue, it would be much more difficult for smaller companies, with far fewer resources in terms of digital maturity and IT infrastructure, to scale up. However, over the years, the costs of data storage[17] and processing[18] have been sloping downward. Therefore, it is reasonable to expect that the scalability issue will

[12] https://tinyurl.com/ov3ulp8.

[13] https://tinyurl.com/jtb8qsq.

[14] https://tinyurl.com/y9xsq5bh.

[15] https://tinyurl.com/6yrby4.

[16] These larger platforms are also known for biasing the recommendations they make towards their own products. The European Commission has considered this practice to be anticompetitive in relation to Google Shopping services, see European Commission, 27 June 2017, Case 39,740, Google Search (Shopping). See https://tinyurl.com/ycpyrqqb.

[17] https://tinyurl.com/ybaatqjn.

[18] https://tinyurl.com/y7akxk7w.

be resolved in a foreseeable future. However, it should be noted that scaling up is by itself not a sufficient condition to enable firms to better price discriminate via advanced algorithms. The longstanding algorithmic challenges remain a far more important prerequisite to be resolved first.

Another challenge is that most firms lack the appropriate technological infrastructure, such as electronic price tags in retail, required to get insights into customer behavior or predicting the WTP of individual customers (Reinartz et al. 2017). Currently, firms might not have sufficient incentives to invest in such technological infrastructure as there are no guarantees that these investments would yield additional profits, or whether the profits would outweigh the investment cost. However, the advent of RFID[19] technologies should enable most firms to equip themselves with state-of-the-art digital retail and sales infrastructure in the short term. Such an infrastructure would often comprise of components, including electronic price tags and electronic shelf labels. However, as with scalability, the availability of such components and digital infrastructure is by itself insufficient to enable firms to better price discriminate.

Finally, while the assumptions made in the scientific literature (e.g. a unique seller, unlimited inventory of a single product, sequential arrival of consumers) are important to ensure the tractability of the proposed algorithms, it remains to be seen how these assumptions can be relaxed or adequately formalized for implementation in real-world applications.

4.1.2 Economic and legal policymaking against the background of technical challenges

Supposing even that some of the aforementioned fundamental, longstanding challenges could be overcome, personalized pricing is not yet guaranteed, as price discrimination does not necessarily and always represent an economic equilibrium. In a monopolistic environment, information on consumers is valuable for the firm, which can increase its profits by price discrimination (Bergemann et al. 2015). Furthermore, additional information leading to a finer and more granular consumer segmentation leads to higher profits for the firm (Belleflamme and Peitz 2015 chapter 8). This may no longer be true in a competitive context. When detailed information on consumers is available, firms will compete for each and every one of them; each consumer becomes a market in its own right. Access to information, therefore, would intensify competition by making each firm's demand more elastic (see Stole 2003). Eventually, price discrimination may not survive competition. This, for example, is the case in Belleflamme et al. (2019), where firms have information about the same set of consumers. In this case, firms cannot escape the Bertrand paradox and apply marginal cost pricing to all consumers. In differentiated product models, firms can compete with personalized prices. In this case, though, competition is still intensified by the use of personalized pricing (see e.g. Thisse and Vives (1988) for a location model, and Stole (2003) for a general model) compared to uniform pricing,

[19] https://tinyurl.com/ycq5bxse.

although firms may not be able to extract all consumer surplus. When firms have access to different sets of data, many papers show that a price discrimination equilibrium could exist (Belleflamme et al. 2019; Montes et al. 2018).

It therefore follows that the amount and structure of the data available to firms determine the equilibrium prices and profits. This entails that how data are produced, traded and acquired is critical for firms. Retailers should decide whether to acquire data (or processed information), and data acquisition is part of the competitive process. Datasets available from brokers, such as comScore, can be considered standard. Thus, competing firms, which purchase such data, will all be on a level-playing field. Liu and Serfes (2004) show that firms are in a prisoner's dilemma situation about deciding to acquire information at the expense of their profits. Data brokers should also decide which data they sell and to whom. Several papers show that data brokers are better off if they sell exclusive datasets (Montes et al. 2018) or different ones (Belleflamme et al. 2019; Bounie et al. 2018) to consumers. Data brokers can extract more information from retailing firms when they compete with asymmetric data. In such a scenario, it is apparent that the firm having access to the data will be better equipped to implement finer-grained PD. Consequently, we can expect that firms will engage in intensified competition in the data market.

In such a competition context, the susceptibility of algorithms to prediction errors and the inherent stochasticity of machine learning models may decrease competition among firms in a market, as even with the same dataset, different algorithms can produce different consumer segmentations.

Finally, the real-life cases of PD reported in mainstream media have raised consumer awareness of how the personal data that they reveal on online platforms can be exploited in pricing decisions. Consequently, they now act strategically and regulate their online behavior in order to distort the personal data that they disclose or hinder its harvesting (Townley et al. 2017). Such actions render personalized pricing more difficult, even in the context of differentiated products, where consumers can offer different data to different producers or sellers. Consumers may also prevent data from being revealed to firms. By investing in data protection, consumers can limit access to their data. Montes et al. (2018) and Belleflamme and Vergote (2016) have models where information disclosure by consumers is a strategic choice.

The legal framework in place at the European Union level would further stimulate such competition. This framework thus could have an additional chilling effect on the emergence of price discrimination strategies. In the EU, data protection law has been upgraded and updated with the general data protection regulation or GDPR (Regulation 2016/679). Enacted on 25 May 2018, the GDPR at first sight imposes a dense regulatory framework on businesses, which seemingly discourages a true market in data.

As a matter of substance, the GDPR determines that data can only be collected for specified, explicit and legitimate purposes but not further processed in a manner that is incompatible with these purposes. The use of data, besides, is limited to these purposes (art. 6). In essence, the GDPR sets up a regime based on consent. The transfer of data from one operator to another is above all possible when a data subject explicitly asks for this (art. 20). In contrast, transferring data without the subject knowing it, is considered an infringement of the GDPR principles.

Upon closer analysis, the ambition of the GDPR is precisely to ensure the free movement of data across the different European Union Member States. The Regulation even acknowledges this at great length, indicating that the free movement of personal data within the Union shall be neither restricted nor prohibited if all requirements of GDPR have been complied with (art. 1(3)). In practice, it turns out that compliance with the GDPR, as such, would not stop a market for data from flourishing, and this for two reasons. First, the requirement for obtaining consent from data subjects appears to have been satisfied relatively swiftly. It suffices to highlight how rapidly any Google or Facebook user consents to these businesses transferring their personal data by ticking some boxes when logging in or making an account. Although the GDPR states that users must understand that they consent to their data being used, most users do not think twice before giving their consent. Second, GDPR enforcement is entrusted to different national data protection supervisors. In this respect, the sanctions and deterrence mechanisms in place, in particular, target businesses that have no measures allowing individuals access to and control over their data. Once individuals have consented to their data being transferred, the threat of sanctions alone would not effectively stop data controllers and processors from transferring data. Data protection authorities, besides, have limited time and resources, so their priorities remain to be seen: one can only fear that they would be too busy investigating the failure of data processors to ensure individuals' access to their own data to oversee the flow of data. So, again, the mere presence of sanctions does not automatically rule out the persistence or emergence of a market in data.

It follows from the foregoing discussion that, although the GDPR inevitably imposes new obligations on businesses controlling or processing data, it does not as such limit the emergence of a market in data. The emergence of such a market would be pro-competitive and limit the possibility for competing firms to price discriminate as long as the technical challenges outlined above remain unaddressed.

4.2 Tacit collusion: technical challenges and economic and legal policy responses

The literature review presented above (Sect. 3.3) appears to provide scientific and empirical evidence supporting the hypothesis that algorithms can learn to collude tacitly. However, the collusive behavior of most (if not all) algorithms discussed were assessed in strict, laboratory and experimental (lab) settings. Such settings can be considered, at best, an artificial manifestation of the various intricacies actually present in the real world.

As explained below (Sect. 4.2.1), it is unlikely that sophisticated machine learning-based pricing algorithms can be effectively deployed in real-world highly competitive markets, with a large number of differentiated products and low-barriers to entry (in some industry), at least in the near future. This corroborates the findings of other researchers, such as the Competition and Markets Authority (CMA Report 2018) and Schwalbe (2019), which also report that "algorithmic collusion currently seems (far more) difficult to achieve" Schwalbe (2019).

To the extent that these real-life scenarios are unlikely to be implemented, the need for significant modifications to antitrust policy or laws would appear somewhat premature (Sect. 4.2.2).

4.2.1 Technical challenges

Real world markets are characterized by the presence of multiple agents, simultaneously exploring and learning about their environments. Past studies (Tan 1993; Tesauro and Kephart 2002) have shown that classical Q-learning algorithms could be deployed in multi-agent environments. It is well-known, though, that multi-agent settings pose several challenges to Q-learning (Ittoo and Petit 2017; Klein 2018; Lowe et al. 2017; Foerster et al. 2018), as discussed below.

The first issue, which pertains to scalability, is the explosion in space requirements as the number of agents grows. As demonstrated in Ittoo and Petit (2017), for an environment described by m states, consisting of n agents, and where each agent can perform A possible actions, the total space requirements is nmA^n, which is exponential to the number of agents. The experiments of Calvano et al. (2018b) showed that increasing the number of agents from $n=2$ to $n=3$ and $n=4$ led to a drastic explosion of the Q-matrix's size, from 3375 to around 50,000, and over 750,000 entries. However, scalability being tangential to their study, the authors did not discuss how it could affect the deployment of their methods. Scalability issues can lead to such high memory and processing time requirements that they are unreasonable and beyond the reach of most firms.

In addition, the presence of multiple agents makes the environment non-stationary. As an illustration, consider an agent (agent A) exploring its environment to learn about the values of its actions in certain states. This exploration affects the reward signals of other (competitor) agents, affecting how they update the utility of each state-action pair, which in turn, affects Agent A's learning procedure. Each agent's policy (action sequence) changes as training progresses. The environment becomes non-stationary from the perspective of any individual agent in such a way that is not explainable by changes in the agent's own policy. Agents are confronted with a "shoot a moving target" problem (Ittoo and Petit 2017; Klein 2018; Lowe et al. 2017). Some recent studies investigating cooperation and competition between agents have proposed various mechanisms to deal with the non-stationarity issue (Foerster et al. 2018). However, most of this research tends to overlook the issue of non-stationarity as it considers it orthogonal to the game-theoretic focus (e.g. cooperation and competition). Furthermore, the mechanisms proposed to deal with non-stationarity are specific to agents that rely on deep reinforcement learning, such as deep Q-learning (where a deep neural network is used to approximate the Q-values). They are not generalizable to classical Q-learning, which is the paradigm underlying most (if not all) current research on algorithmic collusion.

Non-stationary environment presents learning stability challenges. Specifically, there are *no theoretical guarantees* that Q-learning agents converge to a collusive or competitive equilibrium in a non-stationary environment (Lowe et al. 2017; Ittoo and Petit 2017) consisting of at least two ($n >= 2$) agents. Multiple interacting agents ($n >= 2$) may still converge, such as in Klein (2018) with $n=2$ agents, and

in Calvano et al. (2018a, b) with n = 2. However, even if agents in a non-stationary environment do converge, it is unclear whether they will learn an optimal policy (Calvano et al. 2018b). The absence of any theoretical guarantees makes it hard to verify convergence. Verifying whether the agents have learnt an optimal policy is even more difficult. To circumvent this issue, Calvano et al. (2018a, b) relied on heuristics; for example, convergence is assumed if for each player the optimal strategy does not change for 25,000 consecutive periods. Such heuristics offers an intuitive way to establish convergence. However, it remains specific to the environment being investigated; different heuristics has to be crafted for different problem setting.

Related to convergence is the issue of temporality, i.e. the time required for the algorithm to explore its environment and learn an optimal or even a reasonably good policy. This issue is compounded in multi-agent environments since, as explained above, the space requirements (Q-matrix size) grow exponentially with the number of agents. In Calvano et al. (2018a), agents learnt and converged towards collusive strategies after 70,000 iterations. In real life, this time period corresponded to half a year (Calvano et al. 2018a). In Calvano et al. (2018b), the required number of iterations was even higher at 500 000. In practice, however, these delays might be far too excessive. Firms will have to wait for at least half a year before deciding whether to collude or not. During this time span, exogenous factors (such as market characteristics) might render the knowledge acquired by the algorithmic agents obsolete.

Calvano et al. (2018b) do not consider the long time taken/needed to converge to be problematic if the algorithms relied on were be trained offline. However, it might be extremely difficult to replicate an online environment in an offline setting. The offline training environment is likely to be different from the online testing environment. There are no guarantees that an algorithm that performs well offline will do equally well when deployed online.

The majority of models (Klein et al. 2018, Calvano et al. 2018a, b) undoubtedly provides a solid basis for assessing whether algorithms can learn how to collude. They, however, tend to overlook some important characteristics and uncertainties of real businesses (van Uytsel 2018). One limitation is the assumption that the products of the competing firms are homogenous. Another limitation is that demand is generally assumed to be linear or deterministic (Klein 2018; Calvano et al. 2018a). Calvano et al. (2018b) addressed this limitation by formalizing the aggregate demand as an "independent and identically distributed" (i.i.d) binary random variable, which can take two equally probable values corresponding to positive or negative demand shocks. Thus, their model captures demand stochasticity, albeit partially. Current models assume that firms can only set a restricted range of prices. In Klein (2018), prices are fixed to integers in a predefined range. Calvano et al. (2018b) assume that prices vary in the range of 10% below Bertrand competition to 10% above monopoly. Furthermore, all models of algorithmic collusion presented to date assume that firms compete solely on the basis of prices (Ittoo and Petit 2017; Schwalbe 2019). They overlook other dimensions of competition such as product differentiation. Given the significant progress made in AI and in TC over the past years, one could nevertheless expect these limitations to be overcome in a foreseeable future.

In addition, it is likely that most competitors will employ pricing algorithms with different technological underpinnings. However, current research has predominantly

focused on games involving Q-learning agents. The question of whether collusion is still observed when different machine learning methods are employed remains unaddressed.[20] We believe that established collusion in such cases will be far more complex if the algorithms are based on different learning paradigms, such as kernel-based methods versus neural-based methods versus tree-based methods. As suggested by Harrington (2018), collusion will be more likely as competitors use the same pricing algorithms from the same provider. This, however, requires several constraints to be satisfied. First, competitors would have to rely on the same data (variables) to "train" their algorithms; and second, the algorithms would have to be optimized (parameterized) in the same manner. However, as discussed earlier, machine learning algorithms are stochastic by nature. They incorporate a degree of randomness and often rely on probabilities in their decisions. In addition, they have to generalize over a given domain (dataset) and make a number of assumptions (Goodfellow et al. 2016). Thus, even if the aforementioned constraints are satisfied, the inherent stochasticity may make it difficult to determine when (if) the algorithms will collude even if based on the same machine learning paradigms (e.g. reinforcement learning, Bayesian learning).

4.2.2 Economic and legal policymaking against the background of technical challenges

The use of pricing algorithms can help overcome obstacles to tacit collusion. But as mentioned earlier, given the current state of the art, it is uncertain whether algorithms can actually learn how to collude tacitly (Ittoo and Petit 2017; Schwalbe 2019). With regards to transparency, algorithmic pricing entails that firms develop automatic tools, such as web crawlers, to automatically gather data available from websites, thereby increasing market transparency and predictability. In such a market, competitors can be instantaneously informed of price changes and can respond directly by undercutting or price matching. Miklós-Thal and Tucker (2019) show that AI that can better monitor and forecast market conditions are more likely to destabilize collusion by making deviations from the collusive outcome more profitable. Hence, increased transparency provided by the algorithms is likely to make the market more, not less, competitive. As a matter of policy, insisting on more transparency—as is currently the central tenet of debates—still appears a sufficient way forward.

The issue of potential tacit collusion has also raised questions from a competition law point of view. Scholars have argued that the existing (United States) competition law provisions are not sufficient to prohibit instances of tacit collusive behavior in the absence of a proven intent to collude (see Ezrachi and Stucke 2016). As such, competition law provisions would have to be adapted in order to ensure that instances of tacit collusive behavior are captured and prohibited. Other voices have particularly called for a renewed focus in enforcement practices. These proposals

[20] Brown and MacKay (2019) have shown that the use of superior pricing algorithms generally leads to significant increases in markups, even in the absence of collusion.

seem to assume that the competition law rules are broad enough to capture tacit collusion, yet call for particular compliance or enforcement frameworks to be put in place. One proposal made would be to ensure that enforcement authorities can monitor algorithms to ensure their compliance with competition rules, although this task would be 'gargantuan' (Harrington 2018). Another proposal would allow businesses to correct the input in their algorithms if they engaged in collusive behavior, implicitly requiring those businesses to monitor their algorithms themselves in a continuous way in order to ensure compliance with competition law (Massarotto 2019).

In the European Union, one study particularly held that the concepts of EU competition law have been interpreted broadly enough to allow tacit collusive behavior to be captured (Blockx 2018). The legal framework requiring no further adaptations does not mean that algorithm-based businesses should not be made aware more explicitly of the need to comply with EU competition law (Ezrachi and Stucke 2016). From a legal point of view, raising awareness, rather than modifying the rules in place, is the more likely major challenge.

4.3 Towards technology-aware economic and legal policymaking for algorithmic price discrimination and tacit collusion

The previous sections demonstrated that technical challenges currently still limit the possibilities for algorithmic PD and TC to emerge. These challenges notwithstanding, the use of algorithms, even if based on simpler methods than those presented in the literature, can still enable firms to make profits, albeit marginally, above the competitive level. As regards PD, a more refined PD would enable firms to extract more surplus (at the intensive margin), but also to expand the demand (at the extensive margin). For most firms, the real question is whether the investment needed to implement and deploy these algorithms will outweigh the profits to be reaped. As noted in Calvano et al. (2018a) and as described above (both for PD and TC), the development of algorithms entails substantial experimentation (e.g. until a suitable policy is learnt). Experimentation, however, is costly, as it entails a sacrifice of profits, at least in the short-term (Calvano et al. 2018a). For this reason in particular, we believe that the hype surrounding the capabilities of algorithms and the potential harm that they can cause to societal welfare is currently unjustified. As a result, the current legal frameworks and economic policies in place do not need to be modified fundamentally in the short term.

Nonetheless, our conclusion that current legal and economic policies do not require immediate modification does not mean that policymakers have no role to play. Given the rate at which AI technologies mature and are adopted in ubiquitous applications, legislators (and society in general) need a better understanding of the mechanics of the possibilities afforded by AI in pricing applications. This will lead to a better appreciation of the harm and of the benefits afforded by AI technologies, and enable the crafting of adequate legislations to ensure an ethical use of AI.

4.4 Overall conclusions

As part of this technology-awareness, it is important to understand that the practices of PD and TC develop according to different modalities. First, from a game theory perspective, PD can be the outcome of a static one-shot game,[21] while TC can only emerge as the outcome of a dynamic repeated game. Second, a classic result in game theory is the existence of multiple collusive equilibria in repeated games. Multiplicity and complexity are concerns for the emergence of a collusive outcome, as firms have to coordinate on multiple dimensions (price, punishment period, etc.). There is not always such a concern in PD as a price discrimination equilibrium can coexist with a flat price equilibrium. Third, the price structure fundamentally differs in PD and TC. TC is based on flat observable prices while PD requires personalized prices that are not always as easily observable. Fourth, PD and TC have different data requirements. Specifically, algorithms for PD require data on consumers, including those from which demographics and behavioral information can be extracted. Conversely, TC algorithms rely less on consumer data but generally require access to data about the markets, such as prices and sales volumes. In addition, the sharing of data may either result in more competition making PD less likely or in more differentiation between products, which facilitates PD but removes arbitrage. It will depend on which and how data are being shared to determine the concrete impact of PD practices. In TC, the sharing of data (e.g. competitors' accessing each others' prices and sales volumes) can be an impetus for collusive practices in any case.

From a technical point of view, PD algorithms are based on classical machine learning (including econometrics) methods, while TC is best modeled as reinforcement learning. PD algorithms, in their current manifestations, are mostly instances of supervised learning, where learning is achieved from past annotated data. By contrast, in TC, the algorithm is required to explore its environment and learn from its reward signals. Hence, it can be considered a form of unsupervised learning, notwithstanding the predefined reward signal. This distinction translates into the rate of technological maturity. Technology enabling PD, especially based on supervised learning, can be considered to mature at a fast rate. Our review highlighted that we are still a long way off mature technology, based on reinforcement or deep reinforcement learning, for TC. Table 1 summarizes the differences.

The abovementioned differences in technical focus and use of AI are important starting points for the design of a technology-conscious legal and economic policy. Our analysis showed that, rather than speculating on when and how the law is to be modified, attention to technological developments and technical challenges becomes key when shaping economic or legal policy. The different stages of maturity of different technologies and the developments in the realm of AI make a close watch on how technologies develop ever more necessary. Lawyers and economists would also need a clear understanding of these technologies prior to designing or proposing

[21] PD is often analyzed as a static one-shot game with firms having access to profiling data even though behavior-based PD can be formulated in two-stages; the first stage being the analysis of consumer behavior, and the second, the actual PD execution.

Table 1 Comparison between price discrimination and tacit collusion

	Price discrimination	Tacit collusion
Type of game	Static/dynamic	Dynamic
Type of equilibrium	Possibly involve mixed strategies	Multiple equilibria.
Type of pricing	Personalized prices, coupons or personalized recommendations	Flat prices
Data	Consumer characteristics, taste, preferences and habits	Market data, prices and sales volume
Data sharing	Could make PD less likely under current technical possibilities	Makes TC more likely
Algorithmic underpinning.	ML (incl. econometrics) to infer WTP	Reinforcement learning
Training of AI algorithms	Mainly on past annotated data	Agent learns from rewards when exploring environment. No past annotated data required.
Technology	Fast maturing	Still largely experimental (reinforcement and deep reinforcement learning)

modifications to current models or frameworks. This contribution hopes to have offered a somewhat richer understanding of the technical challenges that need grasping in order to engage in better informed and structured economic and legal policy debates.

References

Adami, C., & Hintze, A. (2013). Evolutionary instability of zero-determinant strategies demonstrates that winning is not everything. *Nature Communications, 4*(2193), 1–8.

Ban, G. Y., & Keskin, N. B. (2017). *Personalized dynamic pricing with machine learning. SSN*. https://doi.org/10.2139/ssrn.2972985. Retrieved January 21, 2019.

Belleflamme, P., Lam, W., & Vergote, W. (2019). *Price discrimination and dispersion under asymmetric profiling of customers*, CESifo Working paper no. 7964. https://ssrn.com/abstract=3498721.

Belleflamme, P., & Peitz, M. (2015). *Industrial organization market and strategies* (2nd ed.). Cambridge UK: Cambridge University Press.

Belleflamme, P., & Vergote, W. (2016). Monopoly price discrimination and privacy: The hidden cost of hiding. *Economic Letters, 149,* 141–144.

Bergemann, D., Brooks, B., & Morris, S. (2015). The limit of price discrimination. *Amercian Economic Review, 105*(3), 921–957.

Blockx, J. (2018). Antitrust in digital markets in the EU: Policing price bots. In J. M. Veenbrink, C. S. Rusu, & A. Looijestein-Clearie (Eds.), *digital markets in the EU* (pp. 75–89). Nijmegen: Wolf Publishers.

Bounie, D., Dubus, A., & Waelbroeck, P. (2018). *Selling strategic information in digital competitive markets*, MimeoTelecom Paris Tech.

Brown, Z. Y., & MacKay, A. (2019) *Competition in pricing algorithms*. Harvard Business School Working Paper, No. 20-067, November 2019.

Cahn, A., Alfeld, S., Barford, P., & Muthukrishnan, S. (2016). An empirical study of web cookies. In: *Proceedings of the 25th international conference on world wide web* (pp. 891–901). International World Wide Web Conferences Steering Committee.

Calvano, E., Calzolari, G., Denicolo, V. & Pastorello, S. (2018a). *Algorithmic pricing: What implications for competition policy?* SSRN. https://doi.org/10.2139/ssrn.3209781. Retrieved January 21, 2019.

Calvano, E., Calzolari, G., Denicolo, V. & Pastorello, S. (2018b). *Artificial Intelligence, algorithmic pricing and collusion. SSRN*. https://doi.org/10.2139/ssrn.3304991. Retrieved January 21, 2019.

Cavallo, A. (2018). *More Amazon effects: Online competition and pricing behaviors*. National Bureau of Economic Research. https://www.nber.org/papers/w25138. Retrieved January 21, 2019.

Chen, X., Owen, Z., Pixton, C., & Simchi-Levi, D. (2015). *A statistical learning approach to personalization in revenue management*. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2579462. Retrieved January 21, 2019.

Colombo, S. (2010). Product differentiation, price discrimination and collusion. *Research in Economics, 64*(1), 18–27.

Crandall, J. W., Oudah, M., Ishowo-Oloko, F., Abdallah, S., Bonnefon, J. F., et al. (2018). Cooperating with machines. *Nature Communications, 9*(233), 1–12.

d'Aspremont C., & Gabszewicz J. J. (1986). On the stability of collusion. In: J. E. Stiglitz, G. F. Mathewson (Eds.), New developments in the analysis of market structure. *International Economic Association Series* (pp 243–264). London: Palgrave Macmillan.

DG Justice. (2018). *EU Publications*. https://publications.europa.eu/en/publication-detail/-/publication/ed9ce056-c2cf-11e8-9424-01aa75ed71a1/language-en/format-PDF. Retrieved January 21, 2019.

Dubé, J. P., & Misra, S. (2018). *Scalable price targeting*. Wharton (UPenn). https://marketing.wharton.upenn.edu/wp-content/uploads/2018/01/01-25-2018-Misra-Sanjog-PAPER-targeted_pricing.pdf. Retrieved January 21, 2019.

Ezrachi, A. & Stucke, M. (2015). *Artificial intelligence & collusion: When computers inhibit competition. SSRN*. https://doi.org/10.2139/ssrn.2591874. Retrieved January 21, 2019.

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

Ezrachi, A. & Stucke, M. (2017). *Algorithmic collusion: Problems and counter-measures*. OECD. https://www.oecd.org/officialdocuments/publicdisplaydocumentpdf/?cote=DAF/COMP/WD%282017%2925&docLanguage=En. Retrieved March 25, 2019.

Fischer, C., & Normann, H. T. (2019). Collusion and bargaining in asymmetric Cournot duopoly—An experiment. *European Economic Review, 111,* 360–379.

Foerster, J., Chen, R. Y., Al-Shedivat, M., Whiteson, S., Abbeel, P., & Mordatch, I. (2018). Learning with opponent-learning awareness. In *Proceedings of the 17th international conference on autonomous agents and multiagent systems* (pp. 122–130). International Foundation for Autonomous Agents and Multiagent Systems.

Fonseca, M., & Normann, H. T. (2012). Explicit vs. tacit collusion—The impact of communication in oligopoly experiments. *European Economic Review, 56*(8), 1759–1772.

Garces-Tholon, E., Never, D., & Seabright, P. (2009). The ups and downs of the theory of collective dominance: Using game theory for merger policy. In B. Lyons (Ed.), *Cases in European competition policy: The economic analysis*. Cambridge: Cambridge University Press.

Ghose, A., & Yang, S. (2009). An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Science, 55*(10), 1605–1622.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT press.

Hannak, A., Soeller, G., Lazer, D., Mislove, A., & Wilson, C. (2014). Measuring price discrimination and steering on e-commerce web sites. In *Proceedings of the 2014 conference on internet measurement conference* (pp. 305–318). ACM.

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

Helfrich, M., & Herweg, F. (2016). Fighting collusion by permitting price discrimination. *Economic Letters, 145,* 148–151.

Huck, S., Normann, H. T., & Oechssler, J. (2004). Two are few and four are many: Number effects in experimental oligopolies. *Journal of Economic Behavior & Organization, 53*(4), 435–446.

Ittoo, A., & Petit, N. (2017). *Algorithmic pricing agents and tacit collusion: A technological perspective*. *SSRN*. https://doi.org/10.2139/ssrn.3046405. Retrieved March 26, 2019.

Izquierdo S. S., & Izquierdo L. R. (2015). The "Win-Continue, Lose-Reverse" Rule in Cournot Oligopolies: Robustness of collusive outcomes. In F. Amblard, F. Miguel, A. Blanchet, B. Gaudou (Eds.), Advances in artificial economics (pp. 33–44). Berlin: Springer.

Klein, T. (2018). *Autonomous algorithmic collusion: Q-learning under sequential pricing*. *SSRN*. https://doi.org/10.2139/ssrn.3195812. Retrieved January 21, 2019.

Lee, S. I., Lee, H., Abbeel, P., & Ng, A. Y. (2006). *Efficient L_1 regularized logistic regression*. AAAI. http://www.aaai.org/Papers/AAAI/2006/AAAI06-064.pdf. Retrieved January 21, 2019.

Leibo, J. Z., Zambaldi, V., Lanctot, M., Marecki, J., & Graepel, T. (2017). Multi-agent reinforcement learning in sequential social dilemmas. In *Proceedings of the 16th conference on autonomous agents and multiagent systems* (pp. 464–473). International Foundation for Autonomous Agents and Multiagent Systems.

Lerer, A., & Peysakhovich, A. (2017). *Maintaining cooperation in complex social dilemmas using deep reinforcement learning*. arXiv. https://arxiv.org/abs/1707.01068. Retrieved January 21, 2019.

Liu, Q., & Serfes, K. (2004). Quality of information and oligopolistic price discrimination. *Journal of Economics and Management Strategy, 13*(4), 671–702.

Liu, Q., & Serves, K. (2007). Market segmentation and collusive behavior. *International Journal of Industrial Organization, 25,* 355–378.

Lopez-Feldman. (2012). *Introduction to contingent valuation using STATA*. MPRA. https://mpra.ub.uni-muenchen.de/41018/2/MPRA_paper_41018.pdf. Retrieved January 1, 2020.

Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, O. P., & Mordatch, I. (2017). Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in neural information processing systems* (pp. 6379–6390).

Competition and Markets Authority Report. (2018). *Pricing algorithms. competition and markets authority (CMA)*. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/746353/Algorithms_econ_report.pdf. Retrieved January 21, 2019.

Massarotto, G. (2019). *From digital to blockchain markets: What role for antitrust and regulation. SSRN*. https://doi.org/10.2139/ssrn.3323420. Retrieved March 27, 2019.

Miklós-Thal, J., & Tucker, C. (2019). Collusion by algorithm: Does better demand prediction facilitate coordination between sellers. *Management Science*. https://doi.org/10.1287/mnsc.2019.3287.

Montes, R., Sand-Zantman, W., & Valletti, T. (2018). The value of personal information in online markets with endogenous privacy. *Management Science* (in press).

Motta, M. (2004). *Competition policy theory and practice*. Cambridge: Cambridge University Press.
Ohlhausen, M. (2017). Should we fear the things that go beep in the night. FTC. https://www.ftc.gov/system/files/documents/public_statements/1220893/ohlhausen_-_concurrences_5-23-17.pdf. Retrieved 26 March 2019.
Peysakhovich, A., & Lerer, A. (2017). Prosocial learning agents solve generalized stag hunts better than selfish ones. In *Proceedings of the 17th international conference on autonomous agents and multiagent systems* (pp. 2043–2044). International Foundation for Autonomous Agents and Multiagent Systems.
Rana, R., & Oliveira, F. S. (2015). Dynamic pricing policies for interdependend perishable products or services using reinforcement learning. *Expert Systems with Applications, 42,* 426–436.
Reinartz, W., Haucap, J., Wiegand, N., & Hunold, M. (2017). *Price differentiation and dispersion in retailing*. IFH-Förderer. https://www.marketing.uni-koeln.de/sites/marketingarea/user_upload/Price_Differentiation_and_Dispersion_in_Retailing_Whitepaper_Einzelseiten_Version_1_ohne_letzte_Seite.pdf. Retrieved January 21, 2019.
Schwalbe, U. (2019). Algorithms, machine learning, and collusion. *Journal of Competition Law & Economics, 14*(4), 568–607.
Shiller, B. R. (2014). First-degree price discrimination using big data. *Federal Trade Commission*. https://www8.gsb.columbia.edu/faculty-research/sites/faculty-research/files/finance/Industrial/Ben%20Shiller%20–%20Nov%202014_0.pdf. Retrieved January 21, 2019.
Stole, L. (2003). Price discrimination and imperfect competition. *Handbook of Industrial Organization, 3,* 34–47.
Suetens, S., & Potters, J. (2007). Bertrand colludes more than Cournot. *Journal of Experimental Economics, 10,* 71–77.
Tan, M. (1993). Multi-agent reinforcement learning: Independent versus cooperative agents. In *Proceedings of the tenth international conference on international conference on machine learning* (ICML'93).
Tesauro, G., & Kephart, J. O. (2002). Pricing in agent economies using multi-agent Q-learning. *Autonomous Agents and Multi-Agent Systems, 5*(3), 289–304.
Thisse, J., & Vives, X. (1988). On the strategic choice of spatial price policy. *American Economic Review, 78*(1), 122–137.
Townley, C., Morrison, E., & Yeung, K. (2017). Big data and personalized price discrimination in EU competition law. *Yearbook of European Law, 36,* 683–748.
Van Cleynenbreugel, P. (2020). Article 101 TFEU's association of undertakings notion and its surprising potential to help distinguish acceptable from unacceptable algorithmic collusion. *Antitrust Bulletin*. https://doi.org/10.1177/0003603X20929116.
Van Uytsel, S. (2018). Artificial intelligence and collusion: A literature overview. In M. Corrales, M. Fenwick, N. Forgo (Eds), *Robotics, AI and the future of law* (pp. 155–182). Singapore: Springer.
Varian, Hal R. (2009). Online ad auctions. *American Economic Review, 99*(2), 430–434.
Waltman, L., & Kaymak, U. (2008). Q-learning agents in a Cournot oligopoly model. *Journal of Economic Dynamics and Control, 32*(10), 3275–3293.
Woodcock, R. (2019). Personalized pricing as monopolization. *Connecticut Law Review* 51: Forthcoming. https://ssrn.com/abstract=2972369.
Zhou, N., Zhang, L., Li, S., & Wang, Z. (2018). *Algorithmic collusion in cournot duopoly market: Evidence from experimental economics*. arXiv. https://arxiv.org/abs/1802.08061. Retrieved January 21, 2019.