When algorithms set prices: winners and losers

Discussion paper

19 June 2017

www.oxera.com

Contents

Summary

The digital revolution has led to a significant growth in companies' ability to capture, store and analyse data about their customers, competitors and the wider world, through faster processors, cloud storage, and advances in machine learning. Increasingly, companies are using this information to develop algorithms that set prices for them.

This discussion paper[1] examines how the automation of pricing through algorithms can affect competitive outcomes in markets, and result in different consumers being charged different amounts for the same good or service.

There has recently been extensive press coverage of the risk that price-setting algorithms, using artificial intelligence (AI), could collude among themselves, to the detriment of consumers. Academics, Ariel Ezrachi and Maurice Stucke, were among the first to point out this risk,[2] and their work has influenced several recent speeches and comments by representatives of competition authorities, including the European Commissioner for Competition, Margrethe Vestager,[3] and the Organisation for Economic Co-operation and Development (OECD).[4]

At the same time, others suggest that the use of algorithms can be efficient and procompetitive, leading to outcomes that benefit consumers through faster adjustments to prevailing market circumstances.

This discussion paper explores these two contrasting positions. While the impact of algorithms that use simple rules or formulae to set prices can be assessed in a relatively straightforward way, it is more difficult to judge the more advanced algorithms. These increasingly use AI to adapt and learn as they experience new situations. The way in which AI-driven algorithms learn is highly complex, and, typically, you can't ask them why they did something. An outsider can't 'reverse engineer' the algorithm.

Why do companies use algorithmic pricing?

Algorithmic pricing has clear efficiency advantages for companies that use it—it can be cost-reducing, revenue-increasing, or both.

In some cases, this form of pricing is central to the existence of the market in the first place: for example, it is hard to imagine how the online advertising market could function at anything close to the scale it does without automated pricing procedures—in this case, based mostly around software implementations of auctions. This matters for consumers, as online advertising is the source that enables many online services to be offered free of charge to consumers.

---

[1] The paper was inspired by a meeting in May 2017 of the Oxera Economics Council, a group of leading European academics and officials of the European Commission. The paper was written by Oxera, and the views expressed in it cannot be attributed to the Council members or the European Commission. Oxera is grateful to the meeting participants for their input on this topic.
[2] Ezrachi, A. and Stucke, M.E. (2016), *Virtual Competition: The Promise and Perils of the Algorithm-Driven Economy*, Harvard University Press.
[3] Vestager, M. (2017), 'Algorithms and competition', Bundeskartellamt 18th Conference on Competition, Berlin, 16 March.
[4] OECD (2017), 'Algorithms and Collusion - Background Note by the Secretariat', 9 June, DAF/COMP(2017).

Algorithmic pricing is likely to occur in markets where:

- the costs to serve consumers differ considerably from consumer to consumer, which can be approximated using observable data (e.g. credit and insurance markets);
- demand fluctuates much more rapidly than supply (e.g. airlines, hotels and ride-sharing);
- the price-setter has a wide range of products to price, and algorithmic approaches bring a significant cost advantage (e.g. consumer retail).

Companies' motivations for using algorithmic pricing in each of these areas are clear. But how does it influence outcomes for consumers through competition, and are there winners and losers?

How does algorithmic pricing affect competition and consumers?

While many commentators have focused on the likely problems with algorithmic pricing, its many positive features have the potential to enhance outcomes for consumers in several ways.

Algorithms allow for faster and more accurate price adjustments, taking into account extensive market information. This should improve the matching of fluctuating demand and supply, which makes markets work more effectively and results in better outcomes for consumers in the form of lower prices and their demands being met—for example, shorter waiting times for a ride during peak times. Algorithms can also substantially reduce the costs of setting and changing prices, and facilitate entry by new suppliers, as they can quickly learn how a market works.

Algorithmic pricing can also intensify competition directly. By speeding up responses between competing suppliers, prices may converge to competitive outcomes more rapidly than they would do otherwise.

On the downside, some approaches to algorithmic pricing may be better at sustaining tacitly coordinated outcomes than when prices are set by humans. In particular, algorithms have increased capacity to monitor price movements in the market, and are faster at reacting to changes. In theory, this could enable algorithms to reach collusive outcomes more frequently. In certain situations, algorithms can independently learn to avoid price wars, as a way of maximising profits over the longer term. This could harm consumers, who would not see lower prices. The degree to which such collusion among algorithms is likely to happen in practice is not yet clear.

Algorithms may also be able to facilitate vertical agreements or collusion through a common vertical agent in the market. For example, if many companies use algorithm software from the same provider, one firm's algorithm could anticipate the reaction of those of others, and hence be able to set higher prices. Similar concerns might apply to a platform such as eBay or Amazon Marketplace, through which companies sell their product, and which is involved in setting prices.

Thus, algorithmic pricing may require new approaches to competition investigations, and possibly even to the legal definition of competition infringements. Algorithms that reach tacitly coordinated outcomes will, by their nature, be difficult to identify and interpret. Competition authorities will need to think not only about the tools used to identify issues, but also what constitutes an illegal act when algorithms interact. Likewise, companies using algorithms will need to review and test their pricing practices from a legal and economic perspective to avoid infringing competition law.

While much of the debate in relation to algorithms has been around competition law, there are also important distributional implications of algorithmic pricing—i.e. different consumers paying different prices for the same product. Regulators in sectors such as financial services and telecoms are increasingly looking into these issues.

What are the distributional implications of algorithmic pricing?

One feature of the digital economy is that, despite the availability of a large amount of data on consumers' characteristics, attitudes and preferences, there is as yet relatively little evidence of widespread personalised pricing, other than in industries that have always relied on customer-specific pricing, such as insurance. Many subscription services (e.g. for music streaming and on-demand video) have flat rates across all (or groups of) customers, even though usage, while predictable, varies considerably across consumers, as do the associated costs.

It has been shown that consumers do not like personalised prices. Early experiments by Amazon in setting variable prices for individuals[5] saw an overwhelmingly negative response, even if some consumers probably paid less as a result. This 'punishment' of companies by consumers for treating them differently may be a good reason why we have not seen mass adoption of personalised pricing across digital markets.

What algorithmic pricing is doing, however, is disrupting large swathes of cross-subsidisation in many markets—especially where the costs to serve different customers can vary significantly, such as in the credit and insurance sectors.

While setting uniform prices for all consumers may seem fair, it could be hiding cross-subsidisation between different groups of consumers, and this cross-subsidisation itself may not always be considered to be fair. For example, it is cheaper for banks to serve consumers who use Internet banking than those who visit their branches, and yet often all consumers face the same charges for their current account, or get the same interest rate.

While cross-subsidisation can be economically inefficient, it can also protect consumers. For example, regulators are often concerned about the prices paid by vulnerable consumers relative to those paid by 'sophisticated' consumers. Preventing price discrimination can be a way to ensure that competition for sophisticated consumers benefits vulnerable consumers as well.

Price discrimination is driven not only by the cost to serve customers, but also by customers' willingness to pay or to switch provider. Algorithmic approaches to pricing may identify and exploit these differences between consumers more effectively than prices set by humans.

In this context, what constitutes a fair price? Fairness is hard to define economically, but notions of fairness do exist. In particular, there is an increasing amount of media attention and academic research on algorithms that are used to assess reoffending risk in the US criminal justice system, and

[5] Streitfeld, D. (2000), 'On the Web, Price Tags Blur', *The Washington Post*, 27 September.

whether these algorithms produce results that are ‘fair’.[6] Many of these ideas translate directly into scenarios where some measure of risk is relevant to cost, such as in the credit or insurance sectors.

Finally, as with competition enforcement above, how can regulators intervene in markets where price discrimination is driven by algorithms? In particular, it is difficult to interpret how an algorithm determines what prices to charge; simply banning the use of a particular observable piece of data will not prevent discrimination unless that observable data is completely independent of all other observable data. For example, in car insurance markets, banning the use of gender as a price determinant has not prevented algorithms from inferring the driver’s gender from other information, such as the size and colour of the car.

Are algorithms good for us?

Algorithms are opening up whole new markets, allowing new entrants to operate in existing markets, and helping consumers get better value for money. But they pose new challenges to policymakers, regulators and competition authorities. Traditional approaches to spotting collusive activity, by incentivising whistle-blowers, are unlikely to work with algorithms. Also, it is unclear what constitutes evidence of collusive activity in an environment where there is no record of pricing decisions, and where algorithms are making autonomous decisions based on public domain information.

It is important to note that the two broad concerns about algorithmic pricing are unlikely to arise simultaneously in any specific market. Markets with characteristics that may make them amenable to collusion tend to be less favourable to personalised pricing. Markets where personalised pricing is prevalent do not easily lend themselves to collusion.

The framework for enforcing competition law needs to adapt to a world of algorithmic pricing. This could include monitoring digital markets in much more automated ways, building test environments where the algorithms of companies under investigation can be examined to see how they react to shocks, and asking companies to consider the distributional effects of their pricing policies.

Algorithms are here to stay. Competition authorities and regulators will need to adapt to this new world. Yet it is also important to keep in mind that algorithms often help to break down barriers to competition and make markets more effective and transparent.

[6] See Lum, K. and Isaac, W. (2016), ‘To predict and serve?’, *Significance*, **13**:5, October, pp. 14–19; and Larson, J., Mattu, S., Kirtchner, L. and Angwin, J. (2016), ‘How we analysed the COMPAS recidivism algorithm’, 23 May.

1 How and where algorithms set prices

An algorithm is a set of rules for completing a task or solving a problem. In some cases, an algorithm might be an explicit solution to a particular problem—e.g. finding the shortest path between two points in a network.[7]

When it comes to price-setting, algorithms are computer programs that set prices in an automated way. This practice is becoming increasingly common in online markets, and is starting to appear in offline markets.

Algorithms are also used to automate other business processes. In some cases there may be interactions between this and the way that algorithms set prices, especially in search and marketing campaigns. This discussion paper looks specifically at algorithms in pricing.

1.1 Types of algorithmic pricing

As companies tend to consider the exact nature of their algorithms to be commercially sensitive, they are usually unwilling to share the mechanisms behind their pricing processes. From the academic literature, publicly available information, and Oxera's understanding of the market, we are able to identify four broad approaches to algorithmic pricing.

- **Heuristic**: software that applies simple rules-based approaches to pricing, contingent on the state of the world (i.e. based on the information the algorithm has about the market at a particular moment in time). For example, automatically matching a competitor's price, or raising prices by 10% when stock is low.
- **Analytical**: software that sets prices according to the state of the world, where the pricing rule uses statistical analysis of historical data and is static from that point onwards.
- **Autonomous**: software that sets prices according to the state of the world, where the underlying algorithm might be initialised with historical data, but continuously evaluates performance and updates itself based on observed outcomes.
- **Auctions**: software implementations of auctions, as widely used in the sale of advertising space online, and in retail auction sites such as eBay.

The fourth one is distinct from the rest. Prices set using approaches 1–3 are rarely fully transparent—they are essentially take-it-or-leave-it offers that are updated regularly. By contrast, in auctions (approach 4), prices are determined using a mechanism known to market participants.

Moreover, in practice, companies use a mix of these approaches. Many of those that use algorithmic pricing do not hand over total control to the algorithms themselves. Humans intervene to correct obvious errors or account for data that the algorithm cannot process, and some form of constraint is often built into the algorithm itself in the form of maximum and minimum prices.

[7] Dijkstra's algorithm, for example, is designed to find the shortest path between points or nodes in a graph, and can be used in navigation applications to find the quickest route.

1.2 Where algorithmic pricing is used

Algorithmic pricing is prevalent in:

- **online retail**: heuristic methods are common here, especially those that match in-stock competitor prices;
- **insurance and credit**: riskiness, evaluated by analysing historical data, is an important determinant of prices for most products;
- **airlines and hotels**: prices are set to maximise revenue from relatively fixed levels of output. Some airlines, in particular, use sophisticated learning algorithms.

Some larger businesses implement company-specific solutions in house, but off-the-shelf tools are increasingly available for integration into other business systems using a 'software as a service' model, as shown in the following examples.

- Feedvisor's pricing tool can be integrated with a variety of common online sales platforms, and is claimed to actively learn how to set prices to maximise profit:

> Feedvisor's self-learning algorithmic repricer analyzes your competitive landscape 24/7. It then uses artificial intelligence and big data techniques to determine the ideal price based on your business goals.
>
> The end result? More profits with less effort.[8]

- PricingPRO, provided by PROS, recommends product prices and encodes simple pricing rules, including in response to competitors' prices:

> PROS market-based insights provide real-time algorithmic price recommendations based on customer buying patterns so companies can ensure speed, precision, transparency and consistency across all sales channels.[9]

- Intelligence Node offers Incompetitor, a competitor intelligence service, and Inoptimizer, a rules-based pricing engine:

> An intelligent competitor price monitoring algorithm ensures that you are always competitively priced but never underpriced without a good reason. What's the point in offering discounts when all your competitors have run out of stock?[10]
>
> Inoptimizer's rules based pricing engine matches your volume, pricing and revenue goals with automated price adjustments. Adjustments can be triggered according to occasion, competitor intel, product, category, season and consumer behavior.[11]

In most cases, the algorithms themselves are relatively opaque—most companies treat them as confidential trade secrets. Where information is publicly available, it tends to be generic, focusing on the input data used rather than the methodology.[12]

---

[8] Feedvisor, 'Amazon Algorithmic Repricer', https://feedvisor.com/amazon-repricer/, accessed June 2017.
[9] PROS, 'Prescriptive Price Guidance to Maximize Value on Every Sale', http://www.pros.com/solutions/price-optimization-software/price-guidance/, accessed June 2017.
[10] Intelligence Node, 'Incompetitor: Keep the Competition in Sight at all Times', http://www.intelligencenode.com/products-incompetitor.php, accessed June 2017.
[11] Intelligence Node, 'Inoptimizer: Be on the leading edge of retail analytics', http://www.intelligencenode.com/products-inoptimizer.php, accessed June 2017.
[12] For example, Uber increases rates when demand is higher than normal, and Airbnb reports the variables that feed into its algorithm, as well as some assumptions and observations about interactions.

Algorithmic pricing is not confined to online. Some bricks-and-mortar retailers have equipped their stores with electronic price tags to enable more rapid price changes. Kohl's, a US retailer, has sales in its stores that last hours rather than days. Some toll roads in Texas adjust prices every five minutes to keep traffic moving over 50mph.[13]

1.3 How does an algorithm learn?

Beyond heuristic methods, algorithmic pricing is typically managed by data scientists using machine learning methods. An area of computer science, machine learning focuses on identifying general algorithms that can be applied to a range of situations. In the underlying methods, there is a significant overlap with statistics and econometrics, but the fields have very different analytical cultures. In particular, whereas economists tend to look at the causal structure and interpretation of a system, machine learning focuses on predictive power.

At the simple end of the spectrum, an algorithm might rely on a regression analysis of past sales in relation to prices and other factors considered relevant to sales. Using a fitted model, the algorithm takes the most recent observations of the control variables to adjust prices to maximise total predicted sales. In this sense, the algorithm has 'learned' how to set prices given previous experience.

At the more complex end of the spectrum, the algorithm might use real-world data to continuously learn how to set prices, treating pricing as a 'reinforcement learning' problem. A full reinforcement learning approach would frame the problem of how to set prices in terms of learning a pricing *policy* that maximises the discounted stream of future profits, assuming that this is the company's objective. The objective function could easily be something else, such as market share or revenue, or a weighted sum of several performance measures.[14]

More formally, denoting an action[15] a, state s, discount rate γ, and reward r, the problem is to learn a policy $\pi = P(a|s)$ to solve:

$$Q^*(a, s) = \max_{\pi} \mathrm{E}\left[\sum \gamma^t r_t \mid s_t = s, a_t = a, \pi\right]$$

This formulation of the problem is known as Q-Learning. Other approaches and frameworks are used in the literature, but variants of Q-Learning are currently the most popular.

In effect, the solution needs to take into account the link between actions and states, and that actions today can affect future rewards. The most sophisticated approaches also allow for the possibility that the relationships change over time—the solution tries to balance experimentation between different strategies with refining a core strategy, as shown in Figure 4.1 below.

[13] See Schumpeter, J. (2016), 'Flexible figures', *The Economist*, 28 January, http://www.economist.com/news/business/21689541-growing-number-companies-are-using-dynamic-pricing-flexible-figures, accessed June 2017.
[14] This is fairly common in practice. For example, in training an agent to play Pong, a simple computer game, the agent is likely to improve more quickly if it is rewarded for returning the ball as well as for winning games.
[15] An action in this context might be price changes, but conceptually it could be any type of action, such as how much stock to buy or what to spend on advertising.

Figure 1.1 Reinforcement learning framework

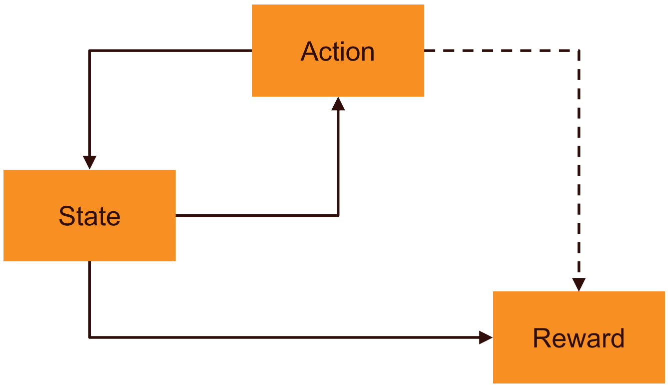

Source: Oxera.

Unlike a simple contemporaneous model of pricing, here:

- the framework acknowledges the circularity in the link between actions (the prices chosen) and the state of the world (competitors' prices, availability, etc.);
- the reward (say, profit) does not have to be provided at the same time as the action is taken.

This type of problem can be solved in various ways, and the field of machine learning is developing rapidly, with significant contributions from technology research labs.[16] State-of-the-art methodologies include variants of 'deep learning', which use artificial neural networks to approximate the optimal action–value function.

The pricing policies that these agents learn are rational, but are driven by unstructured real-world data, and are constrained by the computing power that is available to run the analysis.

[16] This is the same framework that is driving a lot of current AI research, including innovations such as self-driving cars.

2 What economic problems are algorithms solving?

2.1 Benefits of using algorithms

If you are a seller on Amazon Marketplace, whether you use algorithmic pricing is a good indicator of performance on the platform, according to a recent study by Chen et al. (2016).[17] The authors found that users of algorithmic pricing were active for twice as long on the platform, focused on fewer products, and had more positive feedback than non-algorithmic sellers.[18] Algorithmic pricing was also linked with more feedback, suggesting higher sales volumes, which might enable sellers to rank higher in the top sellers list than competitors, including Amazon itself.[19]

Significantly, algorithmic sellers tend to be more successful than their non-algorithmic counterparts at winning the 'Buy Box' for a particular product. The Buy Box is the prominent display of a seller's offer on the product homepage, and, in effect, is the default purchase option through which the majority of purchases are made.[20] This is important because the Buy Box often does not offer the lowest price available within Amazon Marketplace,[21] so the winner might benefit from both higher sales volumes and higher prices than its competitors, and therefore higher profits.[22]

Following the Amazon example, online companies are likely to adopt algorithmic pricing because it increases profit by raising revenues or lowering costs. However, the reason why they might be able to achieve higher revenues or lower costs can differ from one industry to another.

From a consumer perspective, prices might change much more quickly. In particular, when algorithms reduce costs, they might lead to lower prices for consumers, and might help to match supply and demand in the most efficient way. However, faster price changes could also increase search costs, as consumers become less sure about when and where they can expect to find a good deal.

Many platforms that offer matching services, such as Airbnb and Uber, also offer pricing tools. Potential interactions between the matching algorithms and the pricing algorithms can lead to different outcomes than if the platform were performing a matching function alone. Examples of business models that use algorithmic pricing in different ways are given below. In each one, the economic rationale is slightly different:

- insurance companies: compensating for risks;
- Amazon sellers: price-matching;
- airlines: managing perishable goods;
- Uber: adjusting capacity in the short run;

[17] Chen, L., Mislove, A. and Wilson, C. (2016), 'An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace', *Proceedings of the 25th International Conference on World Wide Web* (WWW '16), pp. 1339–49, April.
[18] After identifying sellers who use algorithmic pricing, the authors compared the characteristics of sellers who use algorithms with those of sellers who do not.
[19] Amazon still appears in the Buy Box more often than any other seller.
[20] To select other sellers' offers, the customer has to click a smaller entry in a list of options that appears below the Buy Box listing, or click through to another page to view all entries. The Buy Box winner is chosen by Amazon's own algorithm.
[21] At least 60% of algorithmic sellers did not have the lowest price. See Chen, et al. (2016), op. cit.
[22] Algorithmic sellers win the Buy Box due to their feedback and sales volume, despite asking prices that tend to be higher than those of their competitors.

- Airbnb: pricing a diverse range of products;
- ad placement platforms: encouraging truthful bids.

2.2 General insurance products: compensating for risks

General insurance products, such as motor and home insurance, are typically priced according to the risk characteristics of the policyholder (and the insured item, such as a car). This risk-based pricing allows companies to recoup their costs where an insured event is more likely to arise, and can make insurance markets work more effectively.[23] In practice, this means that insurance products are individually priced, as the risk characteristics (car type, past claims, location, security measures, etc.) are different for each individual.

From using algorithms to set prices for insurance products according to risk characteristics, it is a small step to setting them according to consumers' willingness to pay. This is often done by offering lower prices to new customers (switchers) and then steadily increasing prices for those who remain with the insurer over time. This practice is common, and reflects the highly competitive nature of price-comparison websites,[24] a channel that in several countries is increasingly used by insurance providers to sell their policies. In contrast, many remaining customers are less price-sensitive, reflecting the relatively low levels of engagement that are common to many financial services products.

Some insurance brokers and providers use more sophisticated pricing techniques than simply discounting the first year, often aided by specialist pricing software.[25] The extensive personal data that these companies need to collect to estimate risk can often also be used to estimate willingness to pay. For example, indicators such as employment status and car type can provide clues about whether someone is likely to shop around.

In practice, the extent to which these algorithms can charge more to customers depends on the rates of switching in insurance products—for example, switching rates for motor insurance in the UK are around 40% per annum.[26] In some countries switching motor insurance provider is relatively easy, and is facilitated by the use of price-comparison websites or digital comparison tools. Elsewhere, the take-up of such websites has been much less.

2.3 Amazon sellers on Marketplace: price matching

In a dynamic environment with multiple competitors and quick price adjustments, monitoring competitors' prices can be costly and slow for sellers on Amazon Marketplace if it is not automated. To encourage switching from customers, sellers may use algorithms in order to compete. Users of algorithms represent around 2–3% of total sellers on Marketplace, and around 40% of those sellers that change their product price at least 20 times over its lifespan.[27] These predominantly price-matching algorithms (e.g. matching the lowest or second-lowest price offered for the product) follow a heuristic approach.

[23] For a discussion, see Oxera (2010), 'The use of gender in insurance pricing: unfair discrimination?', *Agenda*, September.
[24] For a description of the dynamic nature of price-comparison websites, see Competition and Markets Authority (2017), 'Digital comparison tools market study: Update paper', 28 March, accessed June 2017.
[25] One example is Earnix, accessed June 2017.
[26] See Competition and Markets Authority (2014), 'Private motor insurance market investigation: final report', 24 September, para. 8.21.
[27] Chen et al. (2016), op. cit.

Taking the prices set by sellers of a particular product, Amazon uses its Buy Box algorithm (see Figure 2.1) to rank merchants by 'best' offering. At an individual level, sellers have an incentive to lower their prices in order to be featured in the Buy Box, but Amazon's Buy Box algorithm also accounts for other dimensions of the sellers' offering that Amazon deems relevant—for example, how positive the seller's feedback is, the number of reviews that the seller has received, and the identity of the seller, especially whether it is Amazon itself.

Figure 2.1 The Amazon Buy Box

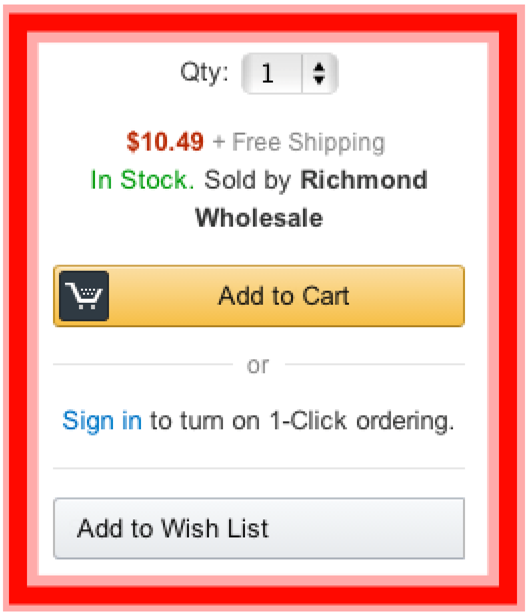

Source: Amazon.

2.4 Airlines: managing perishable goods

Airlines usually face low short-run marginal costs and high fixed costs. They can also face intense competition in some parts of their network, especially for certain fare categories on the most popular routes. In addition, the services they sell are 'perishable': if a ticket is not sold and the plane has flown, the potential revenues are lost forever.

The commercial challenge is that, once capacity has been allocated, airlines must fill the planes at the best price they can. Fluctuations in demand make this difficult. If prices are set too high, planes may end up with empty seats; if too low, planes may be full, but the airline would have forgone potential revenues. To tackle this issue, airlines begin selling tickets long before the flight date, and use sophisticated analytics to dynamically adjust prices up to that date in order to maximise revenue.[28] Car rental companies and hotels use similar approaches to maximise revenue from their relatively fixed levels of capacity, although, to Oxera's knowledge, their approaches tend to be less sophisticated than those used by airlines.

[28] For example, easyJet uses an algorithm that mostly sets prices autonomously. See presentation to Innovation Enterprise by Alberto Rey-Villaverde, Head of Data Science for EasyJet, 'easyJet: data science & innovation'.

2.5 Uber: adjusting capacity in the short run

Uber offers a ride-hailing app that allows customers to book a car at short notice. The customer's smartphone can indicate the pick-up. Once a customer and a driver are matched, an estimated time of arrival is given, and contacts are exchanged. In that way, the consumer can track the driver as they arrive and can communicate with them if needed. After the ride, payment occurs through the app automatically, and the customer and driver rate one another. This two-way rating between customers and drivers means that a customer can refuse a ride, as can the driver.

Prices for the ride are set by Uber's algorithm and cannot be influenced by the driver. The final price for the ride varies by distance and time travelled.[29] When there is a shortage of supply (indicated by rising waiting times for a pick-up), the rate increases by a multiplying factor (which Uber calls 'surge' pricing), and returns to the base rate when waiting times begin to fall.[30]

As such, Uber's pricing algorithm aims to continually balance supply and demand in the short run. As the price rises, more price-sensitive customers stop requesting rides, reducing demand, and the higher prices attract drivers to areas where the surge is active. The base rate acts as a price floor: prices do not fall further even if demand is extremely low—the surge multiplier can never fall below 1.

Uber's exact pricing structure varies depending on the local regulatory environment. In general, drivers keep the booking price, less a percentage that Uber keeps as a service fee. In this sense, Uber's incentive as the price-setter is to maximise revenue for the service.

Originally, periods of surge pricing were clearly announced to customers, who had to accept the surge multiplier before placing an order on the app. See Figure 2.2 for an example.

Figure 2.2 Surge pricing display on the Uber app

Source: Uber.

For consumers, there would not be a surprise when the final bill was presented; for drivers, some rides were more profitable than others. In that sense, price transparency appeared to be welfare-improving. However, for

[29] See Uber Help, https://help.uber.com/h/33ed4293-383c-4d73-a610-d171d3aa5a78, accessed June 2017.
[30] See Uber Help, https://help.uber.com/h/e9375d5e-917b-4bc5-8142-23b89a440eec, accessed June 2017.

several cities Uber tested different price framing and stopped announcing prominently when the surge price was in place and when it ended.[31] This may have been, in part, a reaction to surge pricing transparency creating lumpiness in the market: some consumers would delay their journey, while drivers would decide not to take customers, in an attempt to charge higher prices once surge pricing kicked in.[32] These behaviours possibly created dynamics of under- and over-capacity in the market, resulting in inefficiencies.

Some customers saw surge pricing as taking advantage of situations of high demand, such as terrorist attacks, while others understood it as a way to disincentive non-essential travel and encourage drivers to operate when demand outweighs supply.

2.6 Airbnb: pricing a diverse range of products

Airbnb is a marketplace that connects guests with local hosts. Hosts can rent out available space by listing details on the platform and setting their own rent and check-in/-out times. They have control over bookings and can accept or reject requests after reading reviews about the customers. Airbnb deals with advertising,[33] contracting and insurance, so the administrative burden on hosts is relatively low. Guests search for a property via the Airbnb website, and may filter the listings by size, price, amenities, location, etc., depending on their preferences. Payments are centralised, as guests pay through the Airbnb portal, and Airbnb pays the hosts.

Hosts can set prices freely, but Airbnb recommends prices to them according to an algorithm that incorporates machine learning. The price recommendations are based on criteria such as location, the property's occupancy rate, the booking duration, the size of the accommodation, the time of year, and competitors' prices and availability. The recommendations vary over time—for example, to take into account local events—and are updated regularly.[34]

Airbnb charges a percentage of the total booking price to both the host and the guest. The percentage is smaller for hosts and non-linear for guests—i.e. for higher booking prices, Airbnb charges guests a lower percentage. The platform's strategy is therefore threefold: (i) to maximise the number of transactions; (ii) to ensure that listings are optimally priced; (iii) to ensure participation on the platform by both hosts and guests.

To maximise the number of transactions, guests' search costs and hosts' administration costs must be minimal. In addition, to secure hosts the greatest revenues, and guests the best value for money, prices should be set according to the listings' quality. It is therefore essential that hosts consider prices charged by others in order to ensure a smooth price–quality trade-off across listings. Airbnb's pricing strategy aims to maximise the value of bookings by ensuring that prices are optimal to both parties, and by providing enough incentives for hosts to list their available space on the platform.

[31] See Carson, B. (2016), 'Uber will stop showing the surge price that it charges for rides', *Business Insider*, June, accessed June 2017.
[32] Hood, J.R. (2016), 'Uber drivers cancel rides at last minute, consumers complain', CONSUMERAFFAIRS, 29 December.
[33] Airbnb also offers photo services via its network of freelance photographers so that listings can benefit from high-definition pictures.
[34] See Yee, H. and Ifrach, B. (2015), 'Aerosolve: Machine learning for humans', Airbnb, 4 June, accessed June 2017.

2.7 Ad placement platforms: encouraging truthful bids

Another use of algorithmic pricing relates to ad placement. The mechanisms used allow Google or Yahoo to place ads in search results. The algorithms also determine the price for a sponsored link by eliciting (at least partially) advertisers' willingness to pay.

In the case of Google AdWords, the service is based on two components: the information about viewers that Adwords collects via cookies, and the keywords that advertisers specify.

When viewers search for a keyword, AdWords selects the corresponding ad using personal data, and displays it. When users divert their browsing to click on the ad, advertisers pay a fee. The fees[35] are based on the number of clicks, the number of impressions, or the number of conversions.[36]

To choose the right links to sponsor, AdWords asks companies to bid for keywords. Each time a viewer searches on Google, AdWords runs an auction to determine which ads show on the search results page.[37] The price that advertisers pay is calculated as the minimum amount required to exceed the rank of the next ranked ad.[38]

This mechanism is similar in spirit to a Vickrey auction: the ad placement is awarded to the highest bidder, paying the second-highest price. However, the economics literature highlights that these algorithms, referred to as 'generalised second price' (GSP) auctions, differ from a standard Vickrey auction, in that there are several slots to allocate (i.e. for each keyword there may be several sponsored links).

[35] See Google AdWords, Help page, accessed June 2017.
[36] Impressions: the number of times an ad is shown in a viewable position; conversions: the number of times viewers take a specific action after clicking on one ad leading to a website.
[37] When placing a bid, advertisers can use a traffic estimator to see how many clicks they are likely to get, choosing an 'exact' or a 'broad' match. Google then ranks the offer by bid, weighted by the bid's predicted click-through-rate. It puts the best ads at the top of the page on the basis of the bid's ranking.
[38] See, for instance, Varian, H. 'Search engine advertising', accessed June 2017.

3 How does algorithmic pricing affect competition?

Various commentators have pointed to the possibility that pricing algorithms could facilitate existing forms of collusion and even give rise to new forms of collusion.[39] This contrasts with the suggestion that the use of algorithms can be pro-competitive and lead to outcomes that benefit consumers more quickly and make it harder to maintain collusion. These issues are explored below, assessing the competitive implications of pricing algorithms more widely.

We first consider the likely competitive benefits of pricing algorithms, and then turn to the potential for coordination involving competitors and different parts of a value chain.

One aspect not explicitly covered in the discussion below is that algorithms could also become a dimension of competition, with companies aiming to develop algorithms that are 'better'—for example, by adjusting prices more quickly, by better anticipating changes in demand, or by better matching prices to consumer preferences. This could be pro-competitive if consumers benefit from the increased efficiency in the market. However, in certain markets, such a tendency could also favour market concentration towards the company with the most successful algorithm.

Academic research on the economic impact of algorithmic pricing is as yet relatively limited. Much of the existing research focuses on large platform operators, especially Uber and Amazon. The current literature may not be definitive, however, and there is room for significant empirical evidence to test the competing theories on the impact of algorithms.

3.1 Competitive benefits

Companies can define algorithms to price their goods for a variety of reasons, many of which are also likely to benefit consumers, as follows.

Faster price adjustments

Algorithms can be faster and better at correctly identifying changing market conditions such as demand shocks and cost changes. This enables companies to adjust prices more quickly to the efficient price level. This, in turn, reduces instances of excess supply and excess demand, especially where there are capacity constraints, thereby increasing market efficiency. For example, if companies compete on price and face a cost reduction, they could adjust their prices and reach a new competitive equilibrium much more quickly. This could reduce frictions—for example, in the form of strong demand fluctuations for companies with marginally different prices.

Algorithms used by platforms set prices that clear markets at every state of the world.[40] Many platforms structure their businesses to align their incentives with maximising output on the platform, especially where platforms hold detailed information about demand and supply, and one or both of these changes frequently. In these cases, the platform operator has an informational advantage over the individual sides of the platform.

In this sense, a platform can use an algorithm to set prices that are much closer to competitive prices than they would be in a scenario where users set

[39] See, for example, Lynch, D.J. (2017), 'Policing the digital cartels. Price-setting algorithms mean regulators must now tackle collusion among machines', *Financial Times*, 8 January.
[40] This is the role of the 'Walrasian auctioneer', who receives information about each person's demand at every given price of a good and sets the price such that total demand equals total supply.

the prices. As discussed in section 2, Uber's surge pricing incentivises drivers to offer a ride in times of high demand, even if opportunity costs are high, such as on New Year's Eve.[41] Airbnb also adjusts its recommended price based on demand in order to balance it with supply—for example, when there are local events.

Cost reduction

Algorithms can monitor the market and adjust prices at a very low marginal cost. Limited human involvement reduces staff costs and may reduce the scope for behavioural biases (such as people's tendency to prefer avoiding a loss to acquiring a gain of the same magnitude, referred to as 'loss aversion'). However, at the start, setting up the algorithm and verifying that it is behaving 'well' can be costly. Smaller providers can reduce this upfront cost by buying a software subscription, such as Feedvisor. These long-term cost reductions may eventually be passed on to consumers in the form of lower prices.

Lower barriers to entry

For retailers, algorithms (and subscriptions to algorithm software) could reduce the amount of specific market knowledge required to enter a market. Alternatively, existing retailers might find it easier to broaden their product offering and include products about which they may have less expertise. In this case, the algorithm would set the price for a larger group of products.[42] This could increase the number of companies offering individual products, which would lead to stronger competition.

Algorithmic pricing is likely to achieve some of the listed benefits when used by any kind of company. Some new business models (such as Uber and Airbnb) rely on algorithmic pricing as an essential part of their strategy. In these cases, it is difficult to disentangle the pro-competitive effects of the business model as a whole from that of the pricing algorithm.

3.2 Potential harm to competition

The European Commission and various national competition authorities have started to review pricing algorithms more closely.[43] Scenarios of how algorithms could facilitate anticompetitive outcomes are set out below, before we consider in more detail the theories of harm for horizontal and vertical coordination.

3.2.1 The Ezrachi and Stucke taxonomy

Ezrachi and Stucke (2016)[44] develop various theories of harm for algorithmic pricing. They classify algorithms according to their role in the price-setting process, assess their potential in facilitating collusion, and explore the legal 'toolbox' available to competition authorities in such situations, as follows.

- **Messenger**: to help the monitoring and maintenance of a cartel, which in itself would be prosecuted as an illegal agreement.

---

[41] Hall, J., Kendrick C. and Nosko, C. (2015), 'The Effects of Uber's Surge Pricing: A Case Study', accessed June 2017.
[42] Current evidence suggests that it is more common for retailers who use algorithmic pricing to have narrower product ranges. See Chen et al. (2016). However, there are no insights yet into the direction of causality, whether these retailers may otherwise not have entered at all, and whether this might change in the future.
[43] Vestager (2017), op. cit.
[44] Ezrachi, A. and Stucke, M.E. (2016), *Virtual Competition: The Promise and Perils of the Algorithm-Driven Economy*, Harvard University Press.

- **Hub & Spoke**: companies could agree to use a single price-setting algorithm to coordinate on collusive prices.
- **Predictable Agent**: each company could unilaterally choose an algorithm that monitors its rivals' behaviour, punishes deviations from the collusive price, or otherwise facilitates tacit collusion.
- **Autonomous Machine**: potentially unaware that collusion might be a likely consequence, companies could employ a learning pricing algorithm that is set a target such as profit maximisation. With self-learning, a problem of parallelism—with similar data and similar goals—might arise, and algorithms could find it optimal to increase transparency and collude on higher prices.

The companies' intention is important for the legal assessment of these scenarios—i.e. the Messenger scenario amounts to explicit collusion and is therefore clearly illegal under existing competition rules. In contrast, the Predictable Agent and Autonomous Machine scenarios are less clear because intent is not required to reach collusive outcomes. In these cases, algorithms can facilitate tacit collusion, which the current legal framework does not capture directly. In the context of economic analysis, however, intent is not necessarily the decisive factor—tacit collusion can lead to outcomes that are as bad as explicit collusion.

From an enforcement perspective, it is important to strike a balance between correcting clearly anticompetitive market outcomes and attributing any deviation from perfect competition to parallel behaviour induced by algorithms. While legal tools for tacit collusion are currently limited, a better understanding of the effects of algorithms can help focus on areas where harm is most likely to occur in practice.

3.2.2 Direct coordination

One of the most prominent cases to date where companies agreed to fix prices and used pricing algorithms to implement and control the agreement is one where retailers sold posters on Amazon Marketplace in the USA.[45] The retailers used commercially available algorithm pricing software to coordinate prices. This software operated 'by collecting competitor pricing information for a specific product sold on Amazon Marketplace and applying pricing rules set by the seller'.[46] The UK Competition and Markets Authority (CMA) penalised two retailers for similar coordination in the UK that also relied on pricing software for Amazon Marketplace.[47]

It is not clear whether, in the absence of such explicit agreements, algorithms facilitate coordination. This relates to Ezrachi and Stucke's theories on the Predictable Agent and the Autonomous Machine.

Algorithms as price matching revisited?

The existing literature on price-matching guarantees (PMGs) provides insights into the analysis of algorithms. PMGs commit a company to a certain reaction to a competitor's price, just as an algorithm might 'automatically' adjust a price in response to a change in the competitive environment.

In a simple setting, PMGs can facilitate collusion by providing a credible threat of direct punishment if the competitor undercuts the collusive price that makes

[45] See Department of Justice (2015), 'Former E-Commerce Executive Charged with Price Fixing in the Antitrust Division's First Online Marketplace Prosecution', press release, 6 April.
[46] See ibid.
[47] Competition and Markets Authority (2016), 'Online sales of posters and frames', Case 50223, 12 August.

undercutting unattractive in the first place.[48] Similarly, such punishment could be implemented as an algorithmic rule and executed in real time. The concerns about algorithms facilitating collusion include the following.

- **Increased transparency**: under PMGs, consumers themselves fulfil the role of 'monitoring' prices of competitors of the company they buy from. Algorithms could collect competitors' prices much more quickly and might also be better at monitoring prices of many companies at the same time.
- **Immediate punishment**: with transparency, companies can match a competitor's price immediately. If this price is the result of coordination, this matching constitutes immediate punishment. This means that deviating from an anticompetitive agreement does not translate into any significant profits because the period during which the price of the deviating company is lower might be extremely short before other companies adjust their prices.
- **Reduced risk of detection by competition authorities**: as the incentive to deviate decreases, an agreement is likely to become more stable. This translates into a smaller risk of the agreement being discovered by a competition authority or ended in another way.

Do algorithms commit to coordination?

There is limited insight so far into whether algorithms do have similar effects to PMGs. Salcedo (2015) explicitly considers the risk of algorithms colluding on prices.[49] He argues that the adjustment of algorithms needs time, and for this reason, companies can commit to algorithms until they can be revised again, and hence can commit to employing the tit-for-tat pricing strategy above until revision is possible. In the case of autonomous algorithms, the way in which algorithms learn could take the role of commitment.

- **Commitment power**: companies can credibly commit to an algorithmic reaction function as long as the algorithm cannot be revised. Time-consuming human programming is necessary in altering a deterministic algorithm. In the terminology of Ezrachi and Stucke, this is the Predictable Agent. The Autonomous Machine, on the other hand, learns and revises itself based on observed market outcomes. Its revision is therefore constrained by computing time and the occurrence of new information.

Salcedo makes a strong claim that, under algorithmic pricing, collusion becomes inevitable. This result hinges on two key assumptions: that companies can fully observe, or 'decode', each other's algorithms (or that autonomous algorithms can completely decode each other); and that companies imperfectly commit to an algorithm and can revise it over time.[50] According to Salcedo, in equilibrium this leads to a mutual incentive to choose high prices because the timeframe until the next revision opportunity guarantees benefits from collusion. Figure 3.1 presents a possible evolution of algorithms and prices according to Salcedo (2015).

---

[48] See Hay, G.A. (1982), 'Oligopoly, Shared Monopoly, and Antitrust Law', *Cornell Law Faculty Publications*, Paper 1124; Salop, S. (1986), 'Practices that (Credibly) Facilitate Oligopoly Coordination', in J. Stiglitz and F. Mathewson (eds.), *New Developments in the Analysis of Market Structure*, MIT Press.
[49] Salcedo, B. (2015), 'Pricing Algorithms and Tacit Collusion', Working Paper, November.
[50] This assumption removes symmetric algorithms of 'always setting the Bertrand price' from the set of equilibrium outcomes. If, instead, algorithms could not be revised, the setting by both companies of algorithms of 'always Bertrand' forms an equilibrium: given that my competitor plays and commits to 'always Bertrand' with no possibility of changing this in the future, the best response is to do the same. Yet, if my competitor can adjust their algorithm in the future, I can improve by telling my algorithm to play the Bertrand price now, for example—thereby best responding now—and match any price change (larger than the Bertrand price) of my competitor in the future.

Figure 3.1 Example of algorithmic price collusion following Salcedo (2015)

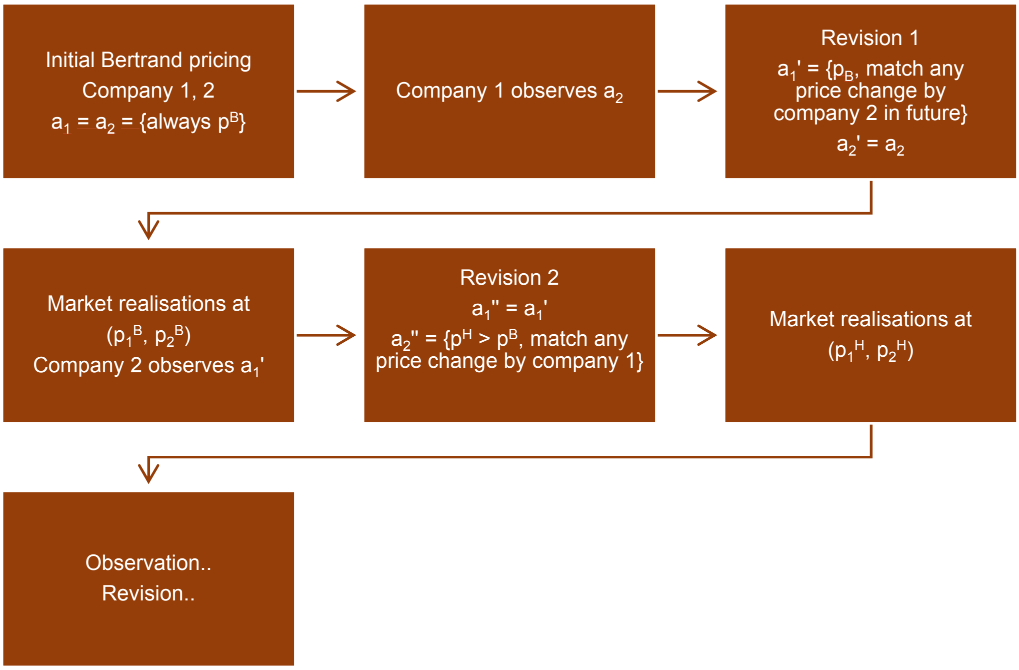

Source: Oxera.

In actual markets, the first assumption is unlikely to be fulfilled. Companies will often use private information or even proprietary technology in the design of their algorithms. This, combined with fast development cycles, makes it seem unlikely that one algorithm would ever be able to decode another from observed outcomes and public information alone. Thus, in Figure 3.1, company 2's algorithm cannot easily 'observe' company 1's 'proposal to collude' (a_1').

Companies might want to circumvent this issue and increase the 'transparency' of their algorithms by revealing their main features. However, such efforts may well be prohibited under current competition rules on the exchange of strategic information between competitors.

Box 3.1 below discusses how machine learning influences the risk of coordination.

Box 3.1 Do experiments in machine learning tell us anything about the risk of coordination?

Autonomous (or self-learning) machines may generate insights into the technical feasibility of ‘colluding computers’. Researchers from Google’s artificial intelligence business, DeepMind, considered the potential for agents built using modern AI techniques to cooperate with one another in a range of settings.[51] Agents were programmed to learn from experience, in the sense explained in section 1, and then to play simple games against, or with, one another. In particular, the research looked at how varying the cognitive capacity of the agents and the parameters of the games changed the outcomes from the games and the strategies used by the agents.

The authors’ main conclusion was that, when the agents are given more cognitive capacity, they appear to be able to sustain more complex cooperative equilibria. But the authors also observed that the agents’ incentive to coordinate was still defined by the environment as well. For example, in a game where agents compete for scarce resources, more cognitive capacity resulted in less cooperation between agents. These findings lead to two insights.

- **Complexity of cooperation** may reduce or prevent tacitly collusive outcomes, even if autonomous algorithms replace humans in the pricing process and cooperation is mutually beneficial. Complexity of collusive strategies will typically increase with the number of competitors, the degree of asymmetric information, the heterogeneity of goals, the complexity and visibility of demand and supply shocks, or if algorithms come from various sources—i.e. if companies develop their algorithms individually.
- **Cooperation is not always rational** and autonomous algorithms might learn to compete vigorously against one another for ever smaller benefits, rather than tacitly cooperate. This is likely to be domain-dependent, and driven by the same features that currently define highly competitive markets.

3.2.3 Vertical coordination

Algorithms may be used not only by companies that want to quickly adjust their prices to changes in the competitive environment, but also by those interested in the prices charged downstream or set on a platform where the good is sold. Three key scenarios and their potential concerns are as follows:

- manufacturers monitor downstream prices, which could facilitate the implementation and enforcement of retail price maintenance (RPM);
- platforms monitor prices set on competing platforms, which could facilitate the implementation and enforcement of most-favoured nation clauses (MFNs);
- platforms set prices on behalf of the sellers or assist them in doing so, thereby facilitating hub-and-spoke agreements.

[51] Leibo, J.Z., Zambaldi, V., Lanctot, M., Marecki, J. and Graepel, T. (2017), ‘Multi-agent Reinforcement Learning in Sequential Social Dilemmas’, Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AA-MAS 2017), São Paulo, Brazil.

RPM and MFN

RPM is generally prohibited under competition law. MFNs have increasingly come under scrutiny by competition authorities in online markets. From an economics perspective, both these practices have efficiency benefits, but might also produce anticompetitive effects.

The European Commission is investigating a case relating to manufacturers of electronic goods that may have limited the retailers' ability to set prices online.[52] According to the Commission, this might have been aggravated by pricing software that reacts to prices set by leading retailers. In a recent submission to the OECD, the Commission notes that algorithms might aggravate RPM practices because compliance by one retailer might lead the algorithms of others to follow.[53]

It seems intuitive that pricing algorithms are likely to increase transparency, not only for consumers and competitors, but also for other market participants. This can improve market functioning, but might also facilitate market monitoring on the supply side, in particular of any agreement between manufacturers and retailers or agents.

The anticompetitive effects of such agreements will increase with their effectiveness, which in turn increases as the cost of monitoring falls. Similarly, to the extent that algorithms could constitute a competitive advantage favouring greater market concentration, vertical constraints are more likely to have a negative effect on competition. However, there is little evidence so far on the extent to which algorithms lower monitoring costs or favour market concentration.

Another potential effect is that pricing algorithms could increase the flexibility of RPM prices. The manufacturer/wholesaler could react to (local) shocks to demand and supply.[54] This would reduce the benefits of deviating from the recommended price and stabilise the arrangement.

Hub-and-spoke

In theory, algorithms may help to coordinate prices between suppliers on a platform, or encourage stronger competition. Similar to platforms, providers of pricing algorithm software (e.g. Feedvisor), if chosen simultaneously by many competitors in one market, might also facilitate collusion or be pro-competitive by coordinating supply and demand more efficiently.

The degree to which platforms can enforce or encourage prices above the competitive level, or to which suppliers can use the platform to achieve this, will depend on the competitive context. It is important to distinguish between concerns that relate to the platform (e.g. network effects leading to a tendency to concentration) and those that are driven by algorithmic pricing. In practice, however, they are likely to be interrelated: algorithms may be less problematic in markets with less concentration and may be less likely to facilitate anticompetitive behaviour when the platform allows for alternative price-setting mechanisms.

---

[52] See European Commission (2017), 'Antitrust: Commission opens three investigations into suspected anticompetitive practices in e-commerce', Brussels, February.
[53] European Commission DG Competition (2017), 'Algorithms and Collusion', 14 June.
[54] Compare to Jullien, B. and Rey, P. (2007), 'Resale Price Maintenance and Collusion', *The RAND Journal of Economics*, **38**:4, pp. 983–1001. The authors note that if retailers face (local) demand shocks, an inflexible RPM price increases retailers' benefits from deviating from the agreement.

When assessing whether algorithmic pricing can facilitate anticompetitive conduct, two important factors are (i) competition between different pricing mechanisms; and (ii) wider constraints on market power.[55] The second factor also captures the strength of alternative offerings—for example, retail stores for Amazon, hotels for Airbnb, and taxis and Lyft for Uber.

Competing pricing mechanisms

A platform makes a choice about how the different sides agree on a price:

- Amazon Marketplace: sellers set a price, often with the help of algorithms;
- eBay: sellers determine a range of parameters (such as whether they prefer an auction or a fixed price);
- Airbnb: an algorithm recommends a price based on the features of the listing and the date;
- Uber: an algorithm determines the price based on the availability of drivers compared to requests.

The platforms trade off the potential efficiencies of algorithms (such as matching supply and demand) with the flexibility given to their users (who may not be able to set their own price). This decision is also likely to depend on the parameters of the specific market, such as the heterogeneity of the product and consumer preferences.

From this perspective, the Airbnb algorithm is less likely to create concerns given that hosts can deviate from the price recommended by the algorithm. If Airbnb's algorithm suggested too high a price, sellers might prefer to set a lower price to avoid lower occupancy rates for their listing.

Sellers on eBay may also find it harder to use algorithms to engage in parallel pricing, given the different price formats. Such concerns would be more pronounced in the context of Amazon Marketplace, where sellers could more easily use the same algorithm to apply the same rules for price-setting.

In contrast, on Uber, drivers have less flexibility as the algorithm price is binding. However, from a competition perspective, it is ambiguous whether, and if so how, drivers should be able to influence prices. Two cases illustrate this:

- in a 2014 case, drivers collectively negotiated higher rates with Uber,[56] indicating that the prices increased potentially above those optimal from Uber's perspective;
- Uber faces a lawsuit by a passenger claiming its algorithmic pricing is anticompetitive, in that it fixes prices among drivers by both a vertical and a horizontal restraint—i.e. in the form of an agreement between Uber and the individual drivers, as well as between the drivers. According to the court, 'drivers agreed to use the pricing algorithm "with the clear understanding that all other Uber drivers are agreeing to charge the same fares."'[57]

[55] Relevant factors include the strength of network effects, multi-homing (using more than one platform), barriers to entry and dynamic competition.
[56] See Katz, E. (2016), 'Uber Algorithm Alleged To Constitute Price-Fixing', *New York Law Journal*, **255**:124, June.
[57] United States District Court Southern District of New York (2016), *Meyer v Kalanick*, 15 Civ., Case 1:15-cv-09796-JSR, Document 37, 31 March.

Hence, prices might be higher if drivers can influence prices in a collective way. One alternative is individual price-setting, which may risk giving up the informational benefit of the platform. Recently, a federal judge in the USA put a hold on Seattle allowing Uber and Lyft drivers to unionise.[58] Paul (2016) argues that not allowing Uber drivers to unionise but allowing the platform to fix prices among the drivers is inconsistent.[59] From a competition perspective, however, it is possible for the platform to have better incentives to set competitive prices than its users.[60]

In summary, the emergence of algorithms can create large consumer benefits by allowing companies to respond quickly to changing market conditions. At the same time, algorithmic pricing can also amplify competition concerns, in particular where markets are concentrated and conditions are favourable for parallel behaviour.

3.3 Competition, collusion and individualised pricing

While algorithms can adjust prices to react to changes in the market environment, they can also generate prices based on the characteristics of individual consumers. Different market settings favour pricing based on different variables. For example, homogeneous goods (such as books, hotel rooms or household appliances) are more likely to give rise to pricing based on the market environment, while financial services are more likely to be priced on consumers' characteristics. For the latter, competitive concerns are less likely to arise because companies are less likely to engage in parallel pricing if there are many price points to consider.

It is important to note that the two broad concerns about algorithmic pricing are unlikely to arise simultaneously in any specific market. Markets with characteristics that may make them amenable to collusion tend to be less favourable to personalised pricing. Markets where personalised pricing is prevalent do not easily lend themselves to collusion.

In section 4 we look at algorithmic pricing that allows for differential pricing across consumers and its implications.

[58] See Bensinger, G. (2017), 'Federal Judge Puts Hold on Seattle Ordinance Allowing Uber, Lyft Union Vote', *The Wall Street Journal*, April.
[59] Paul, S. (2017), 'Uber as For-Profit Hiring Hall: A Price-Fixing Paradox and its Implications', *Berkeley Journal of Employment and Labor Law*, **38**:1.
[60] This depends on the fee structure. If a platform's revenues vary more with the output generated than with the revenues, it is likely to set a competitive price. The fee structure, however, is also determined by the platform itself.

4 What are the distributional implications of algorithmic pricing?

As noted earlier, price discrimination is driven not only by differences in the cost to serve consumers, but also by differences in their willingness to pay or to switch provider. Algorithmic approaches to pricing may identify and exploit these differences between consumers more effectively than prices set by humans.

Without price discrimination there is generally a degree of cross-subsidisation between consumers (they pay the same price despite differences in the costs of serving them). This can be economically inefficient, but has the advantage that it may protect the more captive consumers who have a higher willingness to pay or lower ability to switch. Algorithmic pricing enhances the scope for price discrimination, and in particular personalised pricing. In economics terms, personalised pricing is a form of 'first-degree price discrimination', where each customer is charged according to their full willingness to pay.

In this context, what constitutes a fair price? While hard to define economically, there are generally accepted notions of fairness. The distributional impacts of algorithmic pricing, and how policymakers can approach these, are discussed below.

4.1 Overview of personalised pricing in markets

Some degree of personalised pricing has existed in many markets even before the rise of algorithmic pricing. Examples include indirect differentiation through personalised discounts for consumer products, and price differentiation in insurance and credit markets.

With the growing availability of consumer data, companies are increasingly able to offer a more personalised price to each potential buyer, to the limit of offering a unique price to each buyer based on individual cost of service (in competitive markets) or willingness to pay (in monopolistic markets). The digital economy, in particular, provides the scope for companies to combine information about:

- consumer browsing history, including past purchases;
- consumer devices—the Internet of things provides information on consumer use of devices such as cars, appliances, smart home systems, and health trackers;
- the location and device used for browsing;
- peer information, such as peer preferences, purchasing habits and browsing history.

This information can be used to model and predict more accurately individual willingness to pay, with the potential to approximate first-degree price discrimination. As more data becomes available, there is eventually enough to determine the cost to serve on an individual basis.

Thus, personalised pricing algorithms may become better able to reflect the true cost of serving each individual, reducing the level of cross-subsidy in the market. For example, in motor insurance markets, there are significant gender differences in risk levels of consumers. Although now legally banned in the EU and the USA, the use of gender as a rating factor would allow the market to determine a price that more accurately reflects the risk drivers of consumers.[61]

4.2 Impact of algorithmic pricing on consumer surplus and producer surplus

In many cases, policymakers may consider differentiated pricing to be beneficial for consumers. In the absence of algorithms, a company's ability to price differentially is constrained by its level of information on, and understanding of, consumer preferences. With the increased availability of information, and using algorithmic approaches to refine that information to gain a better understanding of demand and consumers' preferences, companies could, in theory, offer consumers a greater range of products based on individual price–quality trade-offs. If companies are able to provide products to consumers that reflect their cost of service, this might expand the market to consumers with a low cost of service, who would otherwise be priced out of the market.

However, where policymakers consider cross-subsidisation to be a 'good' outcome, they might actually prefer companies to be constrained in their ability to price cost-reflectively. For example, vulnerable customer groups may be receiving a subsidised service, but regulators in many sectors will support a certain degree of cross-subsidisation of these consumers.[62]

The UK Financial Conduct Authority (FCA) has acknowledged that differential pricing in insurance can be pro-competitive as long as the extent of redistribution is deemed acceptable:

> Such pricing structures are not necessarily evidence of weak competition overall and may indeed involve relatively intense competition to attract new customers, even forcing prices to these customers below cost in some cases. Economic profits earned on the back-book may be competed away by offering lower prices to the front-book in some, although not necessarily in all cases.
>
> Overall, it is important for regulators to examine the specifics of cases of cross-subsidy and to understand the business models leading to it to see whether any intervention is warranted to avoid intervening against pricing which is beneficial for consumers.[63]

Its view is that most consumers benefit overall from the dynamics of price competition, but that there is a risk that certain groups who do not switch providers could end up paying persistently high prices.

Acquisti and Varian (2005) assess outcomes in the case of algorithms using data about purchase history to derive estimates of willingness to pay.[64] Their model incorporates the ability of consumers to hide past behaviour through privacy settings and applications, and other such defensive measures. The authors find that where a monopoly seller can set its prices using prior purchase behaviour, it might still not be optimal for a company to completely price-discriminate (i.e. charge each consumer exactly all of their willingness to pay). In this case, the company must offer some benefit to customers to induce revealed information (i.e. 'personalised enhanced services'), which offsets its ability to make use of the information. However, such a result suggests that there may be distributional implications because of differentiation between

[61] See Oxera (2010), 'The use of gender in insurance pricing: unfair discrimination?', *Agenda*, September.
[62] Social tariff schemes are prevalent in the retail water and electricity sectors.
[63] See Financial Conduct Authority (2016), 'Price discrimination and cross-subsidy in financial services', Occasional paper no.22, September, p. 21.
[64] Acquisti, A. and Varian, H.R. (2005), 'Conditioning prices on purchase history', *Marketing Science*, **24**:3, pp. 367–81.

customers who actively employ personal privacy measures and those who do not.

The findings do not assess whether this is profitable in an oligopolistic market with few companies, in which the individually profit-maximising decisions may not be welfare-enhancing for all companies together. The impact of this depends on several factors, including market characteristics. In a monopoly, price discrimination transfers wealth from consumer welfare to producer welfare, whereas in oligopolies, price discrimination can benefit consumers through increased competition.[65] Additionally, consumer-side algorithms search engines and recommendation services can increase competition among sellers. The overall effect is therefore ambiguous.[66]

Other factors that influence the ultimate impact on consumer surplus include the type of product and relationship with buyers.[67] Research into the application of revenue management systems in intertemporal price discrimination (state-contingent pricing often applied in the transport industry, such as changing prices for air tickets) finds similar results: price discrimination is welfare-enhancing overall, but the impact on consumer surplus can be ambiguous.[68]

Companies may also choose strategies that focus less on static profit maximisation but involve strategically pricing to increase market power at a later stage. Such strategies include using algorithms to distinguish between sophisticated and naive consumers, and screening consumers with a high willingness to pay (known as ‘cream skimming’). Online markets facilitate this process for companies by allowing them to present a list of prices to consumers, or by requiring consumers to volunteer personal information in order to make a purchase.[69]

Partitioning, whereby companies segment their customers according to information on their characteristics and behaviour, allows companies to identify and capture more profitable segments of the market and leave consumers with lower willingness to pay to rivals. This practice can allow the price-discriminating company to accumulate more market power from a profitable segment of the market, and has the potential for consumer harm, particularly if the companies are targeting vulnerable consumers or exploiting consumer behavioural biases.[70]

Partitioning can also use algorithmic price discrimination in an attempt to exclude rivals. For example, algorithms have been used to determine a buyer’s location and approximate physical distance from a competitor’s store.[71]

Chen and Zhang (2009) show that, when considering dynamics over time, price competition is restricted by the offsetting effects of a company’s objective of pricing lower in order to compete for customers, but pricing higher to benefit from customer loyalty and to screen out consumers with a lower willingness to

[65] Organisation for Economic Co-operation and Development (2016), ‘Price Discrimination – background note by the Secretariat’, Directorate for Financial and Enterprise Affairs Competition Committee, DAF/COMP(2016)15, 29–30, November, section 3.2.
[66] Brynjolfsson, E. (2013), ‘Will big data create a personalized pricing Nirvana for retailers?’, *Digitopoly: Competition in the Digital Age*, 2 September.
[67] Bergemann, D., Brooks, B. and Morris, S. (2013), ‘The limits of price discrimination’, *Cowles Foundation Discussion Paper No. 1896R*, Yale University, 2 July.
[68] Dupuis, N., Ivaldi, M. and Pouyet, J. (2015), ‘A welfare assessment of revenue management systems’, Toulouse School of Economics working paper TSE-547, 4 January. Acquisti, A. and Varian, H.R. (2005), ‘Conditioning prices on purchase history’, *Marketing Science*, **24**:3, pp. 367–81.
[69] Organisation for Economic Co-operation and Development (2016), op. cit., section 9.1.3.
[70] Ibid., section 4.1.
[71] Ibid., section 9.1.3.

pay. The conclusion is that, on balance, dynamic targeted pricing can expand the market and improve social welfare.[72] In a market with increased consumer information, it is not clear whether this outcome holds, as companies develop more sophisticated tools to partition consumers.

4.3 Impact of algorithmic pricing on distribution of consumer surplus

Distributional issues are increasingly a concern to policymakers. For example, they may have ethical concerns about algorithms being used to identify and target certain customers for higher prices. It is therefore important to consider how algorithmic pricing could/might interact with the existing processes and incentives for companies to price-discriminate, especially given that, with the greater prevalence of algorithms, this type of practice is likely to become increasingly available to companies.

Consumers themselves may also have concerns about algorithmic and personalised pricing. Even those who might benefit from lower prices may be uncomfortable about the fairness of charging different customers different prices at the same time, and, in particular, the idea that companies vary prices according to willingness to pay. There may be greater discomfort with differentiated pricing when this practice is not clearly linked to differential costs of service.

Consumer attitudes towards personalised pricing are likely to be one of the major reasons why the practice is not more widespread now. For example, Amazon's experimental price randomisation in the early 2000s drew criticism for this reason, although it was claimed that variation was not based on consumer-specific characteristics.[73]

4.3.1 Potential for discrimination in algorithms

In markets such as credit and insurance, prices are already commonly set on a personalised basis, and may lend themselves to the application of algorithms. Where such pricing reflects costs, this can make insurance markets highly efficient. However, it has also raised policy and fairness questions in situations where factors that predict the cost of service are correlated with consumer characteristics such as ethnicity, age and gender. Pricing along these dimensions is illegal, discriminatory or considered unfair.

An example of questionable discriminatory pricing can be found among US mortgage lenders, who have been accused of widespread racial bias in their mortgage lending practice.[74] In online platforms such as Airbnb, the use of user profiles has provided evidence of discrimination in pricing.[75]

With the potential rise of personalised pricing in the digital economy, consumer characteristics may, intentionally or unintentionally, influence pricing for a wider variety of products. For example, by using geographic pricing, a prominent online exam tutoring provider was found to be charging higher prices for students with Asian backgrounds, controlling for income level by region.[76]

---

[72] Chen, Y. and Zhang, Z. L. (2009), 'Dynamic targeted pricing with strategic consumers', *International Journal of Industrial Organisation,* **27**:1, January, pp. 43–50.
[73] See Salkowski J. (2000), 'Amazon.com's Variable Pricing Draws Ire', *Chicago Tribune*, October.
[74] See Savage, C. (2012), 'Wells Fargo Will Settle Mortgage Bias Charges', *The New York Times*, July,
[75] See Edelmann, B. and Luca, M. (2014), 'Digital Discrimination: The Case of Airbnb.com', Harvard Business School, Working Paper 14-054, January.
[76] See Angwin J., Mattu, S. and Larson, J. (2015), 'The Tiger Mom Tax: Asians Are Nearly Twice as Likely to Get a Higher Price from Princeton Review', *ProPublica*, September.

Algorithmic pricing can be modelled on historical pricing decisions; bias and discrimination in how past pricing decisions were made can mean that algorithms perpetuate the same discriminatory pricing behaviour. Accusations of discrimination are more difficult to prove against algorithms as opposed to people because the actual mechanics of any pricing decision are less transparent and the default assumption is that algorithms are free from human bias. This is not necessarily a criticism of the use of algorithms—they may be no worse than the humans that they replaced—but there may be scope for them to be used to address this type of issue rather than just inherit it.

Self-reinforcing cycles

One concern about the use of algorithms in policy decisions which may exacerbate discriminatory impacts is the potential to create of self-reinforcing cycles, as illustrated in Figure 4.1.

Figure 4.1 Bias persistence in algorithms

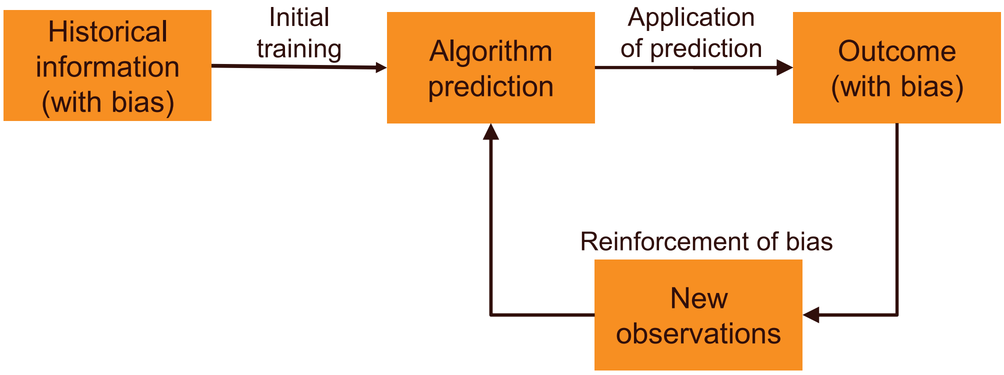

Source: Oxera.

For example, if hiring decisions in the labour market were previously made by human judgement, but done in a discriminatory manner, these decisions can be incorporated into any application-sifting algorithm trained on historical hiring decisions. If this algorithm is then applied to make future hiring decisions, it will continue to do so modelled on previous discriminatory patterns. This makes certain applicant groups more likely to be unfairly rejected, affecting their future employability, and reinforcing the bias. If a wide variety of algorithms is used in the labour market, they may incorporate rejection decisions based on other discriminatory algorithms into their future assessment of applications, creating a self-reinforcing cycle of bias.[77]

This raises concerns because, in markets without algorithms, there may be an active policy objective of reducing bias or protecting vulnerable consumer groups. To the extent that algorithms do contribute to the persistence of these issues, how they do so needs to be understood to enable better algorithm-specific remedies to be developed. Because algorithmic decisions remove the human element of bias, they are more immune to these criticisms (i.e. it can be more difficult to accuse an algorithm of discriminatory intent).

[77] Sometimes referred to as a 'cumulative disadvantage sediment'. See Roth, A. (undated), 'Tradeoffs between fairness and accuracy in machine learning', University of Pennsylvania.

Policy example: PredPol's location prediction tool for predictive policing

Algorithms are used to provide services in the US criminal justice system, from crime prevention to sentencing decisions.

Predictive policing algorithms include PredPol, a location prediction tool used in some metropolitan areas in the USA to predict 'crime hotspots' in order to allocate policing efforts. When comparing the predictions against outturn crime data, Isaac and Lum (2010) found that, rather than focus policing efforts on areas where outturn crime was higher, the predictions focused them on low-income, minority neighbourhoods.[1] Because a crime is more likely to be recorded in an area that is more heavily policed, the outturn data as a result of any algorithmic prediction is self-reinforcing.

Several US states use algorithms in criminal sentencing procedures, including estimating the risk of reoffending (recidivism). The outputs are used to influence decisions to determine bail, probation and parole. Researchers compared two years of ex post data on recidivism rates to the predictions produced by the algorithm COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) for 10,000 criminal defendants. They found the algorithm's results to be a poor predictor of actual recidivism. The predictions from the algorithm were also systematically more likely to mislabel certain racial groups as high-risk, controlling for factors such as criminal history.[2] A similar study conducted on a different predictive instrument, the Post Conviction Risk Assessment (PCRA), found some difference in scoring across racial groups, but found the predictive power of the algorithm to be somewhat more accurate when including a variety of factors, including racial group.[3]

Source: [1] Lum, K. and Isaac, W. (2016), 'To predict and serve?', *Significance*, October, **13**:5, pp. 14–19. [2] Larson, J., Mattu, S., Kirtchner L., and Angwin, J. (2016), 'How we analysed the COMPAS recidivism algorithm', 23 May. [3] Skeem, J.L. and Lowenkamp, C.T. (2016), 'Risk, race, & recidivism: predictive bias and disparate impact', 14 June.

4.4 Regulatory concerns about algorithmic pricing

Evidence on the relative impacts of algorithms on producer versus consumer surplus is currently ambiguous. The impact is also likely to vary from market to market, and needs to be better understood in order to help develop regulatory approaches to address these impacts.

The potential redistribution between consumer groups is best understood using the evidence that algorithms are able to absorb human bias from their training data in public policy contexts. Further work is required to ascertain whether there are any parallel issues in pricing algorithms, if this raises new concerns, and how these should be addressed.

Regulators and policymakers often think about the distribution of outcomes in markets, and actively legislate to reduce the dispersion of outcomes along dimensions that are considered inequitable, such as race or age. If algorithmic pricing introduces distributional effects that are not well understood by regulators—or possibly even the algorithm's creator—this could be problematic: how does a regulator enforce legislation designed to protect vulnerable groups if it cannot evaluate the pricing process that leads to the different outcomes?

4.4.2 Regulating algorithmic bias directly

Examples of direct regulation of pricing include laws prohibiting price discrimination in areas affected by an emergency, or prohibitions against geographic and zonal pricing.[78] For consumer-contingent pricing, there is legal precedent in both the EU and the USA for the ban of price differences based on personal characteristics such as gender, race or disability.[79]

However, it is not clear whether regulating by banning the use of certain inputs will remain effective going forward, as algorithms become more sophisticated and amass a greater level of detailed consumer data. Algorithms might be able to use alternative information to predict 'banned' consumer characteristics. For example, the EU ban on the use of gender in car insurance pricing resulted in a reliance on other information to infer the gender of the applicant. Male drivers typically drive bigger cars or are more likely to work in certain industries. The ban ultimately may have led to a widening of the gap in insurance premiums, as other factors were used to infer the gender of the applicant.[80]

When an algorithm is used to assess risk for a population of users that is composed of subgroups, and the base rate differs between groups, three categories can be used to assess 'fairness' in the predictive power of the algorithm—namely, whether the algorithm:

- makes accurate predictions for each subgroup without systematically assigning higher/lower risk for a subgroup than is actually observed;
- falsely assigns users in one group as high risk relative to another group, across subgroups;
- assigns a higher risk level to individuals within one subgroup relative to another.

Where there are heterogeneous risk levels across subgroups, it can be shown that these three conditions of fairness cannot all hold simultaneously.[81]

In the absence of a direct ban on the use of certain factors, less work is done on how to monitor or detect for bias in algorithms. The most obvious option would be to independently assess whether prediction tools are able to predict outcomes effectively—in the US criminal justice system, the majority of study into the predictive power of algorithms is done by the developers of the instrument themselves.[82]

Algorithmic 'audits' and greater transparency have been suggested as a solution to discrimination in audits, by allowing regulators and the public to understand how an algorithm is making decisions. However, this may not always be feasible, given that the intellectual property behind algorithms will be owned by their creators, and that the algorithms themselves will be considered commercially sensitive.

The design of many algorithms also makes auditing them a challenge, as interpreting the relationships between inputs and outputs from techniques such as artificial neural networks is complex. Transparency of algorithms might also

[78] Organisation for Economic Co-operation and Development (2016), op. cit., section 4.3.
[79] Ibid., section 7.1.
[80] See Collinson, P. (2017), 'How an EU gender equality ruling widened inequality', *The Guardian*, January.
[81] Kleinberg, J., Mullainathan, S. and Raghavan, M. (2016), 'Inherent Trade-Offs in the Fair Determination of Risk Scores', Working paper.
[82] Desmarais, S.L. and Singh, J.P. (2013), 'Risk assessment instruments validated and implemented in correctional settings in the United States', 27 March.

raise competition concerns, as it in effect reveals companies' pricing processes (a theme discussed in section 3).

4.4.3 Assessment of effects

An alternative to regulating algorithmic pricing based on inputs (consumer data) is to focus on the impact (an output-based approach). US regulation prohibits unintentional racial discrimination through the 'disparate impact' doctrine.[83] Traditionally, the enforcement of a disparate impact claim requires a defined group that requires protection, and evidence of a causal connection between the practice and alleged disparities.[84] The application of disparate impact does not require the establishment of intent.

Another option preserves freedom in the competitive process for companies, by allowing the competitive process to rely on price discrimination and algorithms, but with regulators and policymakers considering ex post regulatory interventions in addressing distributional issues by making direct transfers. In the UK energy market, for example, the regulator first intervened by preventing retail suppliers from price-discriminating between incumbent and new customers. However, it later lifted the intervention and focused instead on increasing consumer engagement by mandating lower tariff complexity and standardised presentation of tariffs.[85] This could lead to a similar outcome if consumers who can better understand a market find it easier to avoid price discrimination that is adverse to them.

[83] Organization for Economic Co-operation and Development (2016), op. cit., section 7.1.
[84] See Hancock, P. and Glass, A.C. (2015), 'Symposium: The Supreme Court recognizes but limits disparate impact in its Fair Housing Act decision', SCOTUSblog, June.
[85] See Ofgem (2013), 'The Retail Market Review – Final domestic proposals', 23 April.

5 Concluding thoughts

This discussion paper has highlighted that pricing algorithms could, in many cases, help move markets in favour of the consumer, as consumers themselves use technology to help them find the most competitive prices. The use of algorithms has already opened up new markets, such as free access to online services funded by targeted advertising. Also, by enabling new companies to enter existing markets, such as online travel retail, the visibility of worldwide pricing data increases, helping consumers get better value for money.

However, algorithms will pose new challenges to policymakers, regulators and competition authorities. With algorithms, traditional approaches to finding collusive activity by incentivising whistle-blowers are no longer likely to work. Moreover, in an environment where algorithms are making autonomous decisions based on information in the public domain, and there is no record of pricing decisions, what would constitute evidence of collusive activity is unclear.

When algorithms are used to set prices, it is much easier for them to set prices that vary by customer or group of customer. Evidence from when companies such as Uber and Amazon have tried this type of price discrimination shows that consumers dislike the perceived unfairness of personalised pricing, and can punish companies when they try to use it.

However, flat pricing may not be fair either, particularly when some customers use a service, such as music streaming, much more than others. Algorithms may well start to unwind hidden cross-subsidies and lead to the information we provide about ourselves being used to inform the price we pay. As society starts to look more closely at the distribution of outcomes across people, companies will have to consider whether they are comfortable with the outcomes for different types of consumer before the media or government do it for them.

www.oxera.com

oxera
compelling economics