

Pricing algorithms

Economic working paper on the use of
algorithms to facilitate collusion and
personalised pricing

© Crown copyright 2018

You may reuse this information (not including logos) free of charge in any format or medium, under the terms of the Open Government Licence.

To view this licence, visit www.nationalarchives.gov.uk/doc/open-government-licence/ or write to the Information Policy Team, The National Archives, Kew, London TW9 4DU, or email: psi@nationalarchives.gsi.gov.uk.

Contents

| | <i>Page</i> |
|---|-------------|
| Executive Summary | 3 |
| 1. Introduction | 7 |
| 2. What are pricing algorithms? | 9 |
| 3. The use of pricing algorithms in practice..... | 17 |
| 4. Possible pro-competitive effects of algorithms | 20 |
| 5. Algorithms and coordination | 22 |
| 6. Pricing algorithm simulations | 32 |
| 7. Algorithms and personalised pricing | 36 |
| 8. Features which might raise competition concerns | 47 |
| 9. Further Work..... | 52 |
| Annex 1: Testing for evidence of personalised pricing | 54 |

Executive Summary

1. This economic research paper describes how pricing algorithms are used by firms and explores whether, and under what conditions, the use of pricing algorithms could lead to competition concerns.
2. The paper was prompted by external debate about the role of algorithms in online markets. The main aim was to ensure that the CMA has a good understanding of the existing literature and evidence on the effects of competition on algorithms. We have also carried out a small amount of primary evidence gathering to fill some of the gaps in the literature. We have focused on economic rather than legal analysis.
3. We have found evidence of widespread use of algorithms to set prices particularly on online platforms. For example, many sellers on Amazon use pricing algorithms. As well as simple pricing rules provided by the platforms themselves, some third-party firms sell more sophisticated pricing algorithms to retailers or directly take on the role of pricing using computer models on behalf of their clients.
4. Algorithms can also allow firms to offer different prices to different consumers depending on the information they hold about them. We found limited evidence of this type of personalised pricing in practice, although algorithms are already used to personalize ranking, advertising and perhaps discounts. Increasing data availability, coupled with more sophisticated algorithms, can be expected to increase the scope for firms to engage in personalised pricing in future.
5. There are good reasons to think that the use of pricing algorithms can benefit consumers in many situations. For example, algorithms can reduce transaction costs for firms, reduce frictions in markets, and give consumers greater information on which to base their decisions.
6. A main concern expressed in the literature is that algorithms might facilitate collusive outcomes, leading to consumers paying higher prices. In practice, the concern about collusion covers a broad spectrum of different potential issues. It is important to distinguish between:
 - (a) The use of algorithms to monitor and enforce an existing coordinated strategy; and
 - (b) Theories of harm under which pricing algorithms might lead to coordinated outcomes even when each firm is using the pricing algorithm to make unilateral pricing decisions ('tacit coordination').

7. In relation to facilitating existing coordination, algorithms could make explicitly collusive agreements more stable. For example, algorithms may make it easier to detect and respond to deviations, and reduce the chance of errors or accidental deviations from the collusive agreement. However, the analysis of such agreements is essentially the same as for standard collusive agreements where algorithms are not involved.
8. In relation to tacit coordination, simulation models confirm that some pricing algorithms can lead to collusive outcomes even where firms are each setting prices unilaterally. However, these models typically treat the choice of algorithm as exogenous. This leaves unanswered the question of whether individual firms would have an incentive to deviate, for example by changing the algorithm to undercut the collusive price.
9. Ezechai and Stucke (2016) identify three main ways that pricing algorithms might lead to a tacitly coordinated outcome:¹
 - (a) Hub and spoke – where competing sellers use the same algorithm or data pool to determine prices. This may occur when, instead of using their own data and algorithms, rivals find it more effective to use a third-party algorithm supplier who may gain access to data, or an understanding of their pricing policy from several suppliers. A concern could arise if this gives the platform or ‘hub’ the ability and incentive to increase prices above the competitive level, maximising collective profits.
 - (b) Predictable agent – where pricing algorithms react to events in a predictable way. This could allow the algorithm to signal its intentions and make it easy for competitors to work out what is going on, increasing the likelihood of achieving a tacitly coordinated outcome.
 - (c) Autonomous machine – where algorithms may get so complex and sophisticated that, given an objective to maximise profits, an algorithm can learn by itself and reach a tacitly coordinated outcome, without any intention by its human owners of colluding and limited possibility of discovery by regulators.
10. Of these theories of harm, we consider that the hub-and-spoke concern is likely to present the most immediate risk. This is because it simply requires firms to adopt the same algorithmic pricing model. The predictable agent and

¹ Ezechai, A, and Stucke, ME (2016), *Virtual Competition: The Promise and Perils of the Algorithm-Driven Economy*.

autonomous machine models of collusion could also occur in principle, but rely on pricing algorithms becoming sufficiently advanced and widespread.

11. We have considered the extent to which the possibility of personalised pricing in online markets undermines these theories of harm. In our view, where if there were extensive use of personalised pricing in a market, this would prima facie make it significantly less likely that algorithms could lead to tacit coordination. The traditional conditions that facilitate tacit coordination (such as transparency) make it harder to engage in highly personalised pricing because they mean price comparisons are easier for customers. The increasing use of data and algorithms does not change this. Conversely if pricing is truly personalised then it is difficult for competitors to observe and detect any deviation, making collusion less stable.
12. Finally, we have considered what characteristics of markets or pricing algorithms might make tacit coordination more likely. The main impact of increasing use of data and algorithms appears to be that it can exacerbate traditional risk factors, such as transparency and the speed of price setting. Algorithms can almost instantly observe all competitors' prices, detect any deviation and implement a price response that is objective and easily understandable by competitors.
13. As such, algorithmic pricing may be more likely to facilitate collusion in markets which are already susceptible to coordination, such as where firms' offerings are homogenous. For these 'marginal' markets, the increasing use of data and algorithmic pricing may be the 'last piece of the puzzle' that could allow suppliers to move to a coordinated equilibrium. There could also be greater scope for coordination where algorithmic pricing takes place in an online context where price monitoring and response can happen particularly quickly.
14. One factor which could give competition authorities an indication of whether a price-setting algorithm may result in tacit coordination is the extent to which it leads firms to adopt very simple, transparent, and predictable pricing behaviour (like price matching, or price cycles). Another factor is the prevalence of similar pricing algorithms. If more firms utilise the same pricing algorithm in the same market, it makes it more likely that the market will move to an outcome where prices are higher.
15. Finally, competition authorities could also examine whether the algorithm can place weight on or value future profits. If the algorithm's objective function is very short-term (e.g. maximise profit on each and every sale, with no regard for the impact of its current actions on future profits) then the algorithm is less likely to lead to coordination. For tacit coordination to take place, the algorithm

must be willing to sacrifice short term profits in favour of a longer-term, more profitable outcome. Even for the most sophisticated algorithms that learn to profit maximise over many periods using many variables there should still be a set objective function that the algorithm computes to determine its success. This objective function could in principle be audited by a competition authority, and this may provide some information about the extent to which it is capable of tacit coordination.

1. Introduction

- 1.1 This research project focuses on identifying the conditions under which algorithmic pricing could cause harm to consumers. We have analysed how algorithms might facilitate collusive agreements, how they could result in tacit coordination, and whether there are particular features that make this form of coordination more likely. We have also investigated the use of algorithms to drive personalised pricing, and the interaction between this and collusion.
- 1.2 Algorithms are increasingly used by firms for a wide range of business decisions. This paper focuses on the use of algorithms in firms' pricing decisions, such as setting the market-wide price or offering personalised prices to individual consumers.
- 1.3 Understanding how markets work is an important part of the CMA's duty to promote effective competition for the benefit of consumers. Algorithms and data-based decision making which does not require human involvement are becoming more prevalent. These approaches will continue to develop as access to Big Data and computing power improves.
- 1.4 Algorithms have brought many benefits to consumers and competition in the form of lower costs for suppliers, better service, better product availability, and an improved customer experience. However, it is important for the CMA to understand if and when the use of algorithms might lead to consumer harm.
- 1.5 The analysis in this paper draws on four main sources of evidence:
 - (a) First, we have reviewed the growing competition policy and economic literature which is concerned that the use of algorithms may distort or diminish competition by facilitating explicit collusion or causing tacit coordination;
 - (b) Second, we have contacted a small number of commercial algorithm providers to understand how they operate, and what role pricing algorithms play in market competition.
 - (c) Third, we have spoken to other competition authorities to understand their experience of investigating the use of algorithms.
 - (d) Fourth, we have carried out some pilot tests for the presence of personalised pricing using the CMA internet lab. (See Annex 1 for further detail.)
- 1.6 The aim of the paper is to draw together this evidence and start to identify features of markets and types of algorithms that might raise potential

competition concerns. These factors could help provide a means of preliminary prioritisation for considering future complaints or calls for intervention.

- 1.7 The paper focuses on economic evidence and analysis. It does not seek to assess the lawfulness or otherwise of a given use of pricing algorithms. Nor does it cover broader legal issues such as whether, or in what circumstances, tacit coordination resulting from pricing algorithms could lead to an infringement of competition law.
- 1.8 The remainder of the paper is structured as follows:
 - (a) First, we define what we mean by pricing algorithms, outline how they are currently used by firms, and discuss some of the efficiency benefits flowing from the use of algorithms.
 - (b) Second, we outline the theories of harm by which pricing algorithms might lead to or facilitate collusive outcomes. We describe the conditions under which collusion might be more or less likely to occur. We also summarise some of the literature simulating the possible outcomes of pricing algorithms.
 - (c) Third, we consider the use of algorithms to target personalised pricing offers at individual customers or groups of customers. We focus particularly on the relationship between personalised pricing and the coordination theories of harm, and consider whether coordination and personalised pricing could co-exist in the same market.
 - (d) Fourth, we summarise some of the market features that might raise particular concerns about the use of pricing algorithms either in relation to tacit coordination or personalised pricing.
 - (e) Fifth, we suggest some possible next steps.

2. What are pricing algorithms?

2.1 Algorithms are used for calculation, data processing and automated reasoning. There is not one precise definition of an algorithm that has been universally adopted. Instead there are numerous formal and informal definitions that have been included within the literature. For the purposes of this paper, we adopt the following informal definition:

An algorithm is any well-defined computational procedure that takes some value, or set of values, as input and produces some value, or set of values as output.²

2.2 Algorithms can be specified in English, as a computer program, or even using hardware. The only requirement is that they must be specified using a precise description of the computational steps to be used.

2.3 Algorithms have been developed to solve a wide range of practical applications. This includes algorithms that complete simple tasks such as ordering a series of unordered numbers, to complex algorithms that enable digital encryption, internet communication, and the management of scarce resources.

2.4 The focus of this research paper is on pricing algorithms, and their real and potential effects on how online and offline markets work. Within the broad definition of an algorithm, we define a pricing algorithm as an algorithm that uses price as an input, and/or uses a computational procedure to determine price as an output.

2.5 This definition includes price monitoring algorithms, price recommendation algorithms, and price-setting algorithms. We also consider ranking algorithms, which produce a list of items in an order which is influenced by some training data.

Different levels of sophistication

2.6 Algorithms can be created to address a variety of problems or tasks, from simple to very sophisticated. Pricing algorithms similarly fall on a spectrum of complexity.

² Cormen et al. (2001), [Introduction to Algorithms](#).

Business rules and simple pricing algorithms

- 2.7 Firms have long applied business rules to their operations, including rules on pricing and discounts. Some of these rules can be easily converted into algorithms.
- 2.8 Some pricing algorithms have been designed to follow simple rules such as matching the lowest competitor's price, or remaining within the lowest quartile of prices. For example, Amazon offers a "Match Low Price" feature to third-party sellers on their platform. This allows sellers to match the lowest price offered by competitors, and allows them to choose which competitors to match based on a combination of listing condition, fulfilment method, customer feedback rating, and handling time.³ Automated information collection and pricing could mean that the response to a rival's price change could occur within minutes whereas without an algorithm the response could have taken a few days.
- 2.9 An example is what happened to the price of the book "The Making of a Fly" on Amazon in 2011. This textbook on developmental biology reached a peak price of \$23 million. This price was the result of two sellers' pricing algorithms. The first algorithm automatically set the price of the first seller for 1.27059 times the price of the second seller. The second algorithm automatically set the price of the second seller at 0.9983 times the price of the first seller. This resulted in the price spiralling upwards until one of the sellers spotted the mistake and repriced their offer to \$106.23.⁴ This example appears to have been the result of a lack of "sanity checks" within the algorithms, rather than any anti-competitive intent. However, it demonstrates how the lack of human intervention in algorithmic pricing may lead to unintended results.

Machine learning

- 2.10 Alternatively, a more advanced algorithm could be left to decide what data it considers is most relevant to meeting its objective (such as profit maximising). The algorithm would then act as a "black box" so that even the employees who instruct the algorithm would not know which variables it was using to set a particular price, and may not be aware of whether any increase in profit was due to attracting additional customers, charging higher prices to loyal customers, or tacit coordination. Such complexities may increase when many

³ See Amazon's [Match Low Price Help Page](#).

⁴ This is detailed in a 2011 blog post by Michael Eisen, [Amazon's \\$23,698,655.93 book about flies](#).

of the other firms in the market are also using the same or similar algorithms to set their pricing.

- 2.11 Machine learning algorithms can solve more complex problems. These kinds of algorithms do not have to be explicitly programmed to solve a problem, but can iteratively change and improve by themselves, and therefore they are more flexible than regular hard-coded algorithms.
- 2.12 Machine learning is a very broad field, with many different approaches. A full overview is beyond the scope of this paper. However, it is common to describe machine learning algorithms in terms of three broad categories, depending on the nature of the feedback available to the algorithm:
- (a) Supervised learning – the algorithm is provided with a training set of inputs paired with the correct output ('labels', or 'true' values), and its goal is to work out a function that maps inputs to outputs. The performance of the algorithm can be measured by comparing the predictions of the function when it is applied to a dataset that was not used for training (e.g. cross-validation).
 - (b) Unsupervised learning – the algorithm is provided with data but no labels or examples, and its goal is to find an appropriate function that describes the structure of data (e.g. clustering).
 - (c) Reinforcement learning - in contrast to supervised and unsupervised learning, in which the algorithm is presented with a static dataset (often using historic data, and sometimes referred to as an 'offline' learning), reinforcement learning algorithms iteratively interacts with a dynamic environment. Using the feedback that it gets from the environment, the algorithm tries to work out which actions in the environment will maximise some objective (such as profit).⁵
- 2.13 An early and simple machine learning algorithm developed to set prices is a 'Win-Continue Lose-Reverse' rule, and it commonly serves as a benchmark against which other more sophisticated algorithms are tested. This adaptive algorithm adjusts prices incrementally in one direction and evaluates what

⁵ Reinforcement learning problems often involve trading off 'exploration' (taking potentially suboptimal actions in order to learn about the environment) and 'exploitation' (taking the best action given our current knowledge of the environment). This requires a good algorithm to have some appreciation of the long-term consequences of current actions, even if the short-term reward from current actions is negative – a feature which has parallels with the ability of successful cartellists to be willing to take short-term losses in order to promote or enforce a collusive outcome which is profitable in the long term.

happens to revenue. If revenue increases, it continues to make similar changes to price. If not, it makes an incremental change in the opposite direction. The algorithm make small changes to price in order to learn about market demand, and requires very limited computational resources and no data at all on customers.⁶

- 2.14 Q-learning appears to be a common approach to algorithmic pricing problems. Q-learning attempts to maximises total discounted profit over time, using ‘trial-and-error’ to interact with its environment to learn the optimal pricing policy. It is well suited to pricing because it does not require a model of the environment, such as the demand and competitors’ costs functions. It continuously trades-off between ‘exploiting’ its current knowledge by selecting the action which provides the highest learned payoff, and ‘exploring’ to expand its knowledge by selecting other actions. However, one prominent drawback with Q-learning methods is that it treats the environment as stationary, but the presence of other competitors who are also learning makes the environment non-stationary.

Artificial Neural Networks

- 2.15 One approach to machine learning that has gained recent prominence is Artificial Neural Networks (ANNs), which has been used in applications such as AlphaGo and AlphaZero.^{7,8} ANNs are based on collections of connected units (‘artificial neurons’) that receive, process, and transmit signals to other connected neurons. These connections each have a ‘weight’ that determines

⁶ DiMicco, Greenwald, and Maes (2001), ‘[Dynamic pricing strategies under a finite time horizon](#)’, *Proceedings of the 3rd ACM conference on Electronic Commerce*, pp95-104.


⁷ AlphaGo is a computer program that plays the board game Go, which was developed by Google DeepMind. Go is considered much more difficult for computers to win than other games such as chess. In March 2016, AlphaGo beat Lee Sedol, a 9-dan professional player, without handicaps. Since then, DeepMind has continued developing stronger versions, including versions created without using any data from human games and trained purely by playing games against itself (*tabula rasa* reinforcement learning from self-play). AlphaZero can play Shogi and Chess as well as Go, does not use any input from human games, and in December 2017 it was reported that it beat other world-champion computer programs, including AlphaGo Zero (in Go), Stockfish (in Chess), and Elmo (in Shogi).

⁸ However, it is not clear that the incredible achievements of specialised AI for perfect information games would generalise to solving the problem of achieving stable tacit coordination in real world markets. Although Go has a very large space of possible states, it is still a relatively stable environment with two players, a clear set of rules and permissible actions by players. By contrast, whilst would-be algorithmic coordinators might innovate hard to try and maximise profits within a system of coordination. The maximisation itself can be getting harder as well, especially if rivalry hasn’t been fully suppressed and competitors are innovating equally hard to design adversarial AI that can cheat on a coordinated outcome without detection, if pricing becomes ever more personalised and therefore opaque, or if sophisticated customers innovate on countermeasures. For a discussion of the possible countervailing innovation for customers, see Gal, M, and Elkin-Koren, N (2017), ‘[Algorithmic Consumers](#)’, *Harvard Journal of Law and Technology*, Vol.30.

the strength of signals in the connection. The network 'learns' as the weights are adjusted based on its performance.⁹

- 2.16 The simplest ANNs have a single 'layer' (or level) of neurons, which receives inputs and generates outputs.¹⁰ Deep learning refers to more complex ANNs which have multiple (two or more) layers of neurons, with each layer using the output from the previous layer as input (see Figure 1).

Figure 1: A Deep Learning ANN



Source: Nielsen (2017)

- 2.17 With sufficiently intricate and layered design, deep learning algorithms can be very flexible in their application, and could lead to very nuanced decisions even in complex environments like real-world markets.
- 2.18 Firms can use neural networks to estimate market demand. In contrast with econometric demand models, which usually require some assumption about the functional form of demand, a neural network approach would not need to make any assumptions about demand in advance.
- 2.19 However, deep learning neural networks are also more difficult to understand and it can be hard to tell what is happening in the many layers, creating a

⁹ This is done by minimising some suitably chosen 'cost function', and using methods like gradient descent and backpropagation to overcome the curse of dimensionality created by the high number of weights to be optimised. For our purposes, it is not necessary to explain the details for these methods, but the underlying intuition behind these methods will be familiar to economists who have worked on multivariable optimisation problems. Readers who are interested in these methods may find it helpful to read Nielsen (2017), [Neural Networks and Deep Learning](#).

¹⁰ In many neural network diagrams, inputs and outputs are depicted as nodes and referred to as 'layers'. However, they are not neurons and no processing is involved in the input and output layers.

'black box' so that firms using such networks (and regulators) may not be able to tell the underlying causes of the network's output.

- 2.20 Some companies that sell repricing algorithms claim to use machine learning techniques to improve on simple re-pricing rules. One example of this is an Amazon marketplace algorithmic re-pricer which the CMA contacted (although it is not clear whether they are using a neural network).¹¹ The firm providing pricing services claims to use the Amazon seller's past pricing/profit/revenue data, competing firms' prices, and market information such as competitors' stock levels, to determine the optimal price to charge consumers. Its algorithm also takes into account competitors' publicly-available pricing information and customer feedback. Whereas simple re-pricers often charge the lowest price amongst competitors, this machine learning re-pricer maximises profits through optimising the trade-off between higher prices and lower sales. It adapts to specific business goals such as meeting sales targets, or capturing a specific share of the 'Buy Box' sales (which is the 'default' seller for a product on Amazon).¹²

The relationship between “Big Data” and algorithms

- 2.21 Similar to definitions of 'algorithm', Big Data lacks a universally accepted definition. It is typically defined by three specific characteristics, first described by Laney (2001):¹³
- (a) Data Volume: Big Data is characterised by the large breadth/depth of available data.
 - (b) Data Velocity: Big Data is collected, processed and analysed at a considerably faster speed. Some data is made available in real-time and allows for a process called now-casting (the prediction of the present, the very near future and the very recent past).
 - (c) Data Variety: Big Data typically allows for a wider variety of data to be collected about consumers and their spending habits. As a result, personalised pricing, directed marketing and behavioural discrimination have become a possibility for companies with large datasets.

¹¹ See Annex 1 for further information.

¹² As explained in the Annex 1 on firms selling re-pricers, the Buy Box lists the default retailer for any product listed on Amazon. Most customers do not review the individual sellers and just purchase from the default in the Buy Box if one is listed.

¹³ Laney (2001), *3D Data Management: Controlling Data Volume, Velocity and Variety*.

- 2.22 For most algorithms, data is an essential input. Gal (2017) argues that it is not advances in algorithmic practices that has caused an increase in the prevalence of algorithms, but the dramatic increase in the quality and availability of data. Although there have been significant advancements in efficient algorithms, what has made the difference is the ‘information explosion’ which includes data accumulation, management and analytics.¹⁴
- 2.23 The potential inputs into a pricing algorithm could be any piece of information that would be relevant to price formation, for example:
- (a) competing firms’ prices;
 - (b) firms’ past pricing/profit/revenue data;
 - (c) individual customer information, including their purchase or browsing history or other indicators;
 - (d) market information such as competitors’ stock levels (e.g. whether it is in-stock or not, or more detailed information if this has been made publicly available by competitors trying to communicate scarcity to customers);
 - (e) external information such as weather patterns; or
 - (f) firms’ costs, such as production, storage and fulfilment.
- 2.24 Algorithms can process this information using a set of simple rules, such as price matching the competitor with the lowest price. In this case, the algorithm does not benefit from having past data to draw from. This is because the algorithm does not ‘learn’ from past experiences, but simply chooses prices based on pre-set rules.
- 2.25 A more complex algorithm may rely on a pre-defined prediction model, e.g. regression analysis. In this case, the programmer chooses the factors relevant for sales and the algorithm uses past observations to adjust the model with the aim of maximising revenue. In this case, the algorithm does benefit from having a historical dataset on which to test the parameters of its model to maximise revenue.
- 2.26 Further, algorithms may be able to use real-time Big Data to continuously learn how to set prices using machine learning or deep learning processes (i.e. reinforcement learning). These could allow algorithms to experiment with

¹⁴ Gal, Avigdor (2017), ‘[It’s a Feature, not a Bug: On Learning Algorithms and what they teach us](#)’, OECD Roundtable on Algorithms and Collusion.

different strategies, and to further refine and adjust the pricing model. These algorithms benefit even further from having a historical dataset, as well as rich real-time data, on which to test and update its pricing model.

- 2.27 Although Big Data and algorithms are closely related and discussed together, we focus primarily on the latter in this paper. Therefore, we do not examine whether Big Data, or extensive user data collected and used by incumbents, could be a barrier to entry, create strong network effects, or (in the extreme) whether data could be an 'essential facility'.

3. The use of pricing algorithms in practice

- 3.1 There has been a considerable increase in the sources, types, and volume of data collected by businesses. Ninety percent of the digital data in the world today has been created in the past two years alone.¹⁵ With the growth of Big Data, more businesses are purchasing data analysis services. Worldwide revenue for Big Data and business analytics is predicted to grow from \$130 billion in 2016 to more than \$200 billion in 2020.¹⁶
- 3.2 Pricing algorithms generally fall into two categories:
- (a) Algorithms which are developed by businesses to set the prices for products which they produce and sell to consumers. Generally, they are produced by larger companies with the resources and expertise to develop them.
 - (b) Algorithms which are developed by specialist algorithm development firms. They do not specifically tailor their algorithm to one product or market, and instead license their algorithms for other companies to use. These are sometimes bundled with a broader suite of “business intelligence” services.
- 3.3 Both of these approaches could in principle apply in either online or offline markets, as set out below.

Pricing algorithms in online markets

- 3.4 Pricing algorithms have become prevalent within some online retail markets. Even more than four years ago, in December 2013, Amazon implemented more than 2.5 million price changes every day, a 10-fold increase on the number of price changes in December 2012. This is compared to just 52,956 price changes made by Best Buy and 54,633 changes made by Walmart during November 2013.¹⁷
- 3.5 Not only are large retailers such as Amazon taking advantage of algorithms to re-price their goods, but so are smaller online retailers. A research paper¹⁸

¹⁵ IBM (2016), *10 Key Marketing Trends for 2017*.

¹⁶ IDC (2015), *Worldwide Semiannual Big Data and Analytics Spending Guide*. IDC (2015), *Worldwide Semiannual Big Data and Analytics Spending Guide*.

¹⁷ Profitero (2013), *Profitero Price Intelligence: Amazon makes more than 2.5 million daily price changes*.

¹⁸ Chen et al. (2016), ‘An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace’, *Proceedings of the 25th International Conference on World Wide Web*, pp1339-1349.

developed a methodology to detect whether Amazon third-party sellers are using pricing algorithms to re-price their goods. Analysing the pricing history of the top 1,641 best-selling products, and approximately 30,000 sellers,¹⁹ they predicted that 500 sellers were using algorithmic pricing strategies. These sellers received more feedback and won the Buy Box²⁰ more frequently than non-algorithmic sellers, suggesting higher sales volumes and more revenue. The researchers found that some sellers changed the prices of products tens or even hundreds of times a day. These kinds of price changes would be impossible for a human to replicate, and indicate that algorithms are necessary for sellers to keep pace with the repricing behaviours of their competitors. However, this paper could not conclude whether algorithmic repricing led to higher profits overall, nor how it impacted consumer welfare.

- 3.6 A number of these third-party sellers use algorithmic pricing software developed by specialist algorithm suppliers. In conversations with these sellers, they have indicated that for large Amazon sellers (more than \$1,000,000 in annual revenue), having automatic repricing software is necessary in order to be able to handle their large numbers of products. Further, they have stated that Amazon products can be fiercely competitive markets, and so dynamically responding to changes in competitors' prices is necessary to remain competitive. However, we are not aware of any empirical analysis of what size firms utilise algorithmic re-pricing software.
- 3.7 As part of our research we spoke to two firms that sell pricing algorithm services. One of the firms supplied 'off-the-shelf' pricing algorithms to third party sellers on Amazon Marketplace. The other manages clients' Amazon marketplace pricing on their behalf. We are also aware of other firms that provide bespoke pricing algorithms to individual clients.
- 3.8 The two firms appear to vary in their sophistication, with one offering relatively simple automated business rules and the other claiming to implement more sophisticated machine learning techniques. Both firms aim to help their clients to respond quickly and efficiently to competition and to improve sales. The

¹⁹ Chen et al. (2016) don't state the total number of sellers that they analyse. However, they approximately scrape 20 sellers per product, so we estimate that the total number of sellers that they covered is approximately 30,000. This is likely to be an overestimate, as sellers can sell multiple products.

²⁰ Amazon on the Buy Box: "The Buy Box is the box on a product detail page where customers can begin the purchasing process by adding items to their shopping carts. A key feature of the Amazon website is that multiple sellers can offer the same product. If more than one eligible seller offers a product, they may compete for the Buy Box for that product. To give customers the best possible shopping experience, sellers must meet performance-based requirements to be eligible to compete for Buy Box placement. For many sellers, Buy Box placement can lead to increased sales." (See Amazon's page on [How the Buy Box Works](#).)

firms also outline how it can be useful to understand the pricing strategy of competitors and how price rises at times of the day with low sales volume may reduce competitive pressure.

Pricing algorithms in offline markets

- 3.9 Algorithms are not necessarily constrained to digital or online markets. Algorithms are defined as a procedure that takes inputs and follows a series of well-defined steps to produce an output, and therefore can be implemented without the use of code. Many business rules and processes are simply algorithms. However, the use of data-scraping methods allows for real-time data collection on consumers and competitors which makes it easier to implement algorithmic pricing strategies.
- 3.10 Algorithmic pricing strategies are typically difficult to use in traditional brick-and-mortar retail settings. This is because data, such as competitors' prices, must be collected manually. It is more difficult to collect and store. Re-pricing the products requires manual human intervention in order to physically change the price on offer.
- 3.11 There have been moves to adopt electronic price tags in retail shops. Some major UK retailers, such as Tesco, Sainsbury's and Morrisons are trialling electronic price tags within their shops. These tags allow them to change prices on in-store goods more quickly and frequently in response to fluctuations in demand or to cheaply sell off excess stock.²¹ This would make it easier to implement algorithmic pricing strategies.
- 3.12 Algorithmic pricing strategies have been observed in some offline markets. For example, there have been a number of press reports and academic studies that allege that retail petrol providers have used algorithms to facilitate tacit coordination and improve their profit margins. Although retail petrol is not an 'online' market, there are online websites or services that monitor petrol prices, and individual sites can adjust their prices quickly at almost zero cost.

²¹ Retailer Marks and Spencer trialled a system in 2016 where the price of lunchtime food was made cheaper before 11am, to encourage people to buy their lunch earlier when shops were emptier. See this BBC article from 2017: [Why your bananas could soon cost more in the afternoon](#).

4. Possible pro-competitive effects of algorithms

4.1 In many cases the introduction of algorithms (both pricing algorithms and others) is likely to have positive impacts on consumers and on competition. In this section, we briefly discuss the positive impacts that all algorithms could have on markets by increasing supply-side and demand-side efficiencies. In addition to the direct benefits on the market, algorithms can also assist regulators and competition authorities. One example of this is the cartel screening tool that the CMA has developed to help public bodies and others running procurement.²²

Supply-side efficiencies

4.2 In general, the use of Artificial Intelligence (AI) can significantly reduce labour costs if it is able to replace human workers. For example, one recent survey of machine learning researchers predicts that AI will outperform humans in many activities in the next ten years, such as translating languages (by 2024), driving a truck (by 2027), and working as a surgeon (by 2053).²³ However, AI and algorithms are less likely to be able to perform jobs that require intuition, abstract thinking, or complex physical movements.

4.3 There are further potential efficiencies and cost-savings if algorithms can improve the efficiency of human workers. Mass data collection and algorithmic processing promises to assist managers in making more, faster and better decisions. Where the decision making is also automated, robo-sellers promise still more cost savings.

4.4 More specifically, pricing algorithms may be expected to make markets more efficient and clear faster, as prices become more responsive to changes in supply and demand. Thus perishable goods like groceries or airline tickets are less likely to go to waste, where the remaining stock has no value to the seller but would have some value to buyers. Related to this, pricing algorithms may also enable or facilitate improved inventory management, particularly for perishable goods like hotel rooms and flights.

²² Competition authorities and regulators can and do use algorithms to detect cartels. The CMA has created a cartel screening tool to help procurers screen their tender data for signs of cartel behaviour. This software looks at factors including the text of the bids. It is unlikely the features of collusion relevant to comparing detailed tenders to prevent bid-rigging will be useful in identifying price fixing tacit coordination/collusion in online retail markets.

²³ Grace et al. (2017), *When Will AI Exceed Human Performance? Evidence from AI Experts*.

Demand-side efficiencies

- 4.5 A wide variety of algorithms help consumers make decisions in market transactions. Some offer consumers information that is relevant to their choices. For example, third-party price monitoring tools like CamelCamelCamel help consumers on the Amazon platform purchase products when they are at their lowest price by alerting them when the price for a specific product reaches a certain level.
- 4.6 More sophisticated algorithms use price forecasting to suggest to the consumer whether they should purchase products immediately or wait for an expected decrease in price. An example is the flight price aggregator Kayak, which uses data on previous flight price trends to suggest whether to purchase flights or wait. However, the effectiveness of this algorithm has been shown to be mixed. An article²⁴ showed that for a sample of 15 routes, if a customer were to follow the algorithm's suggestions they would have paid 2% more than if they had instead simply bought the tickets on a random day. It concludes this is likely statistically insignificant.
- 4.7 With a considerable increase in the complexity of algorithms, consumers could completely outsource their purchasing decisions. Digital agents could use data to predict consumer preferences, optimally choose the most suitable product and services, negotiate and execute the transaction, and potentially form consumer coalitions ('buyer groups') to obtain the best terms and conditions. This would significantly reduce search and transaction costs, allow for more sophisticated and rational choices, and even strengthen buyer power.²⁵
- 4.8 Both search and comparison services (in the form of digital comparison tools) and collective buying services (such as Groupon, LivingSocial and Google Offers) already exist. There is speculation that digital personal assistants (like Amazon's Alexa, Apple's Siri and Google's Google Assistant)²⁶ could go beyond what is currently available. For instance, it may be that digital personal assistants can collect and process enough information about users and encourage enough user uptake (e.g. through a better or more seamless user interface), to be able to co-ordinate purchases across a much wider range of users and products than what is currently available.

²⁴ Fung, K (2014), [When to Hold Out For A Lower Airfare](#).

²⁵ Gal, M, and Elkin-Koren, N (2017), '[Algorithmic Consumers](#)', *Harvard Journal of Law and Technology*, Vol.30.

²⁶ IHS Markit (2017), [Digital Assistants to Reach More Than 4 Billion Devices in 2017 as Google Set to Take a Lead](#).

5. Algorithms and coordination

- 5.1 In spite of the benefits of algorithms outlined above, there is a growing competition policy literature raising concerns about the potential for algorithms to lead to consumer harm. One of the main theories of harm relates to the possibility that pricing algorithms might lead to collusive outcomes, with consumers paying higher prices than in a competitive market.
- 5.2 In practice, the concern about collusion covers a broad spectrum of different potential issues. It is important to distinguish between:
- (a) the use of algorithms to monitor and enforce an existing coordinated strategy; and
 - (b) theories of harm under which pricing algorithms might lead to coordinated outcomes even when each firm is using the pricing algorithm to make unilateral pricing decisions.
- 5.3 The section below discusses each of these theories of harm in turn.

The use of algorithms to facilitate explicit agreements

- 5.4 Algorithms may be used as a tool to implement explicit collusion. Below, we give detail on recent cases where algorithms have been used to implement collusive agreements. We then discuss the circumstances in which pricing algorithms could increase the stability of a collusive agreement.


Example of a cartel case involving algorithms

- 5.5 In the CMA's Trod Ltd/GB eye Ltd case,²⁷ the two parties agreed a 'classic' horizontal price-fixing cartel for posters and frames sold on Amazon's UK website. They implemented this agreement using automated repricing software which automatically monitored and adjusted their prices to make sure neither was undercutting the other. The parties kept in contact with each other through regular means to ensure the arrangement was working, and to deal with issues regarding the operation of the re-pricing software.²⁸ Because there was a clear anti-competitive agreement made between humans, the CMA was able to demonstrate that the parties infringed the Chapter I prohibition.

²⁷ Competition and Markets Authority (2016), [CMA issues final decision in online cartel case](#).

²⁸ Note that a similar case was undertaken by the US DoJ; US v Aston and Trod (DoJ, 2015).

Figure 2: Explicit coordination implemented or facilitated by algorithms



Source: CMA

Mechanisms by which algorithms might facilitate collusion

5.6 From an economic perspective, algorithms may make explicitly collusive agreements more stable for a number of reasons:

- (a) it is easier to detect and respond to deviations;
- (b) it reduces the chance of errors or accidental deviations; and
- (c) it reduces agency slack.

Detecting deviations

5.7 Collusive agreements are only stable if the firms can detect when their partners have deviated from the collusive price. Without detection, one of the firms would be able to lower its price, increase its sales and therefore increase its profits. This results in the collusive agreement breaking down.

5.8 If firms are able to detect a deviation, they are then able to 'punish' the deviating firm by lowering their prices even further. This means that the deviating firm is able to enjoy higher profits in the period before 'punishment', but once they have been undercut, their profits fall to below the collusive level. Therefore, the speed at which a deviation is detected and punished influences the incentive for the firms to deviate. The faster the deviation is detected, the lower the expected profits from deviation, and therefore the cartel is more stable.

5.9 Pricing algorithms make the detection of deviation quicker and less costly. This is due to the greater availability of pricing data, both in terms of speed at which it is communicated and the volume of available data. This makes it

easier for competitors, even without the use of algorithms, to monitor prices.²⁹ The likely impact depends to some extent on where in the supply chain the firms operate. At the retail level, it is simple for firms to collect pricing data because prices are typically transparent. With intermediate products, however, prices may not be transparently advertised and therefore pricing algorithms may have little to no effect on how deviations are detected.

Reducing chance of errors

- 5.10 One way in which a cartel can break down is due to ‘noisy price information’. This occurs when firms in the agreement do not receive perfect information about what their co-conspirators are charging. This might then lead, for example, to a seller confusing a period of unusually low demand, and hence low sales, with cheating by its cartel partner.³⁰
- 5.11 Algorithms could make explicitly collusive agreements more stable by making errors such as these less likely. The increased ease and availability of mass data collection makes it easier for firms to accurately understand how their competitors are pricing. Effectively, an algorithm could function similarly to a resale price maintenance agreement, where an upstream cartel is certain that their agreed-upon prices are being followed.³¹

Reducing agency slack

- 5.12 Another feature that reduces the stability of cartels in traditional economic models is “agency slack”. This occurs when, although a collusive agreement has been agreed on between senior managers within a firm, salespeople and other non-management employees may have incentives to undermine the cartel. They may do this if they favour immediate payoffs rather than the long-term benefits of maintaining a cartel; or there may be intra-firm competition for promotions or sales-linked salary rewards. For these reasons, they may choose to undercut the collusive price.³²
- 5.13 Using algorithmic pricing could reduce the possibility that agency slack will lead to the cartel breaking down. This is because there is less scope for

²⁹ Algorithms allow firms to automatically process large volumes of data; which may not be feasible if prices had to be manually monitored. Because these make it quicker for the firm to respond to a deviation, the gain that the firm gets from deviating will be relatively less valuable, and the cartel will therefore be more stable.

³⁰ Green and Porter (1984), ‘[Noncooperative Collusion under Imperfect Price Information](#)’, *Econometrica* Vol.52, pp87-100.

³¹ Mehra, S (2015), ‘[Anitrust and the Robo-Seller: Competition in the Time of Algorithms](#)’, *Minnesota Law Review*, Vol.100 (forthcoming).

³² *ibid.*

individuals within an organisation to take pricing decisions themselves which might go against the collusive agreement.

Could algorithms result in tacit coordination or conscious parallelism?

5.14 In addition to facilitating explicit collusion, some commentators have expressed concerns that pricing algorithms could lead to tacit coordination. This section considers theories of harm under which pricing algorithms might lead to coordinated outcomes even when each firm is using the pricing algorithm to make unilateral pricing decisions.

Alternative theories of harm

5.15 Ezrachi and Stucke (2015)³³ describe three main ways in which algorithms could result in the formation of a tacit coordinated pricing outcome: hub-and-spoke; predictable agent; and autonomous machine.

Hub-and-spoke

5.16 The first way in which algorithms may lead to a tacitly-collusive outcome is when sellers use the same algorithm or data pool to determine price.

5.17 If multiple competitors use the same pricing algorithm, this may lead the competitors to react in a similar way to external events, such as changes in input costs or demand. Furthermore, if the competitors are aware or able to infer that they are using the same or similar pricing algorithms, firms would be better able to predict their competitors' responses to price changes, and this might help firms to better interpret the logic or intention behind competitors' price setting behaviour. Widespread knowledge and use of common pricing algorithms may therefore have a similar effect to information exchange in reducing strategic uncertainty, which may help sustain (but not necessarily lead to) a tacitly coordinated outcome.

5.18 There are some caveats to this theory of harm:


(a) First, the mere fact that firms use the same pricing algorithm is not, by itself, sufficient to establish a tacitly coordinated outcome. There must still be some intention on the part of competitors to acquiesce to the tacit suppression of rivalry. This is because if firms using the same pricing

³³ Ezrachi, Ariel; Stucke, Maurice – [Artificial Intelligence & Collusion: When Computers Inhibit Competition](#) (2015)

algorithm can reach a collusive outcome in which prices are higher than the competitive level, they still need to decide to maintain their strategy and resist the temptation to modify their algorithm or switch strategy to undercut the collusive price and earn higher (short-term) profits. It is not obvious how common pricing algorithms, by themselves, can help firms to overcome this central problem of maintaining collusion, particularly if firms can change or modify their pricing algorithms as easily as they can change individual prices – albeit that algorithms can accelerate the learning process, because price setting is much more rapid.

- (b) Second, how would a firm know that its rivals are using the same algorithms? In a real-world market, many factors could be influencing firms' pricing behaviour, and it may not be possible to fully deduce or reverse engineer a competitors' pricing algorithm. Those seeking to coordinate prices may need to explicitly announce or communicate the details of their pricing algorithm to rivals, which would be akin to explicit coordination.³⁴

Figure 3: 'Tacit' coordination due to common pricing algorithms



Source: CMA

5.19 Arguably a more serious situation is if competitors decide, instead of using their own data and algorithms, that it is more effective to delegate their pricing decisions to a common intermediary which provides algorithmic pricing services. This may result in a hub-and-spoke-like framework emerging, even though competitors are not expressly fixing the price.

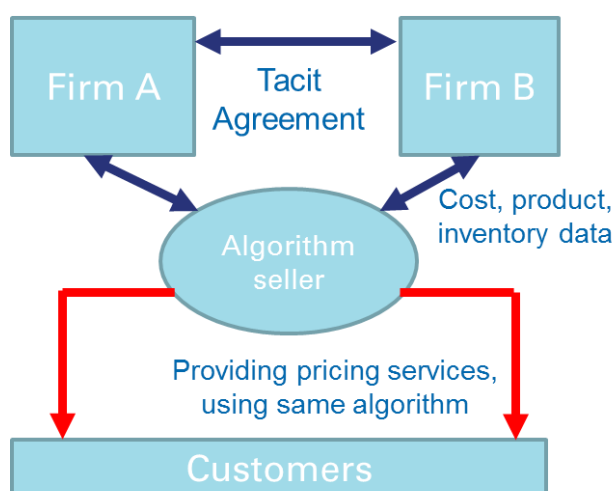
5.20 Relevant considerations for this theory of harm include:

³⁴ As with price transparency more generally, there may be a tension between a consumer law objective of increasing transparency about how firms use pricing algorithms to set prices, particularly if this involves using consumers' data, and an anti-collusion objective of preventing competitors from figuring out their rivals' pricing strategies.

- (a) The proportion of the relevant market that has delegated its pricing to a common intermediary's pricing algorithms.
- (b) Whether the common intermediary's pricing algorithm makes use of non-public information or data from multiple clients (competitors) when determining prices for each client.³⁵
- (c) Whether the objective function of the pricing algorithm is to maximise the total joint profit of all the common intermediary's clients, perhaps because the common intermediary's remuneration is calculated as a proportion of all its clients' sales.

5.21 If a sufficiently large proportion of an industry uses a single algorithm to set prices, this could result in a hub-and-spoke structure that may have the ability and incentive to increase prices. In this scenario, existing competition law analysis of hub-and-spoke could be sufficient to address competition concerns if certain criteria can be established.

Figure 4: 'Tacit' coordination due to common intermediary



Source: CMA

Predictable agent

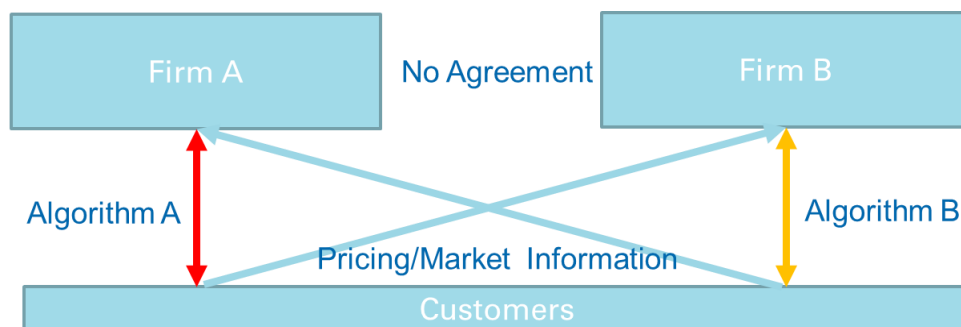
5.22 The second category is that of a predictable agent. Here, humans unilaterally design pricing algorithms which react to external factors in a predictable way. Again, this would have the effect of reducing strategic uncertainty, which may

³⁵ This may not be limited to the inputs used by the algorithm to set/recommend prices. There could still be competition concerns if there was an exchange of historic, competitively sensitive, non-public information during the development (i.e. the 'training' phase) of the algorithm, even if no such data were further supplied during the 'live' phase of the algorithm being used to recommend/set prices.

help sustain (but not necessarily lead to) a tacitly coordinated outcome. The algorithms can be programmed to monitor the market prices, rationally follow price leadership, and punish deviations from a tacit agreement.

- 5.23 In the absence of explicit communication, tacit coordination appears to be more likely to be a concern if the price-setting algorithms leads firms to adopt very simple, transparent, and predictable pricing behaviour (like price matching, or price cycles), which can be recognised by other firms.

Figure 5: Tacit coordination without agreement due to algorithms



Source: CMA

Autonomous machine

- 5.24 The third category is that of the autonomous machine. Here, competitors unilaterally design an algorithm to reach a pre-set target, such as the maximisation of profit. If the algorithm is sufficiently complex, it can learn by itself and experiment with the optimal pricing strategy. There is the possibility that the algorithms may find the optimal strategy is to enhance market transparency and tacitly collude. The important difference with the Predictable Agent model is that the algorithm is not explicitly designed to tacitly collude, but does so itself through self-learning. It is similar to the Predictable Agent model in that it would appear difficult to categorise this as falling within Article 101. The algorithms are not just sustaining existing coordination but generating this coordination themselves.

Reasons why tacit coordination may be more likely as a result of algorithmic pricing

- 5.25 Tacit coordination refers to an anti-competitive market outcome which is achieved without the need for explicit communication between competitors. Below, we consider the reasons why algorithmic pricing may make tacit coordination more likely.

Market Transparency

5.26 A paper by the OECD³⁶ details how the prevalence of algorithmic pricing may result in greater market transparency. It argues that in order for a firm to adopt algorithmic pricing, it must first collect detailed real-time data on its competitors. Therefore, they have an incentive to develop automated methods to collect and store data without human intervention. Once some market players invest in the systems needed to benefit from algorithmic pricing, the remaining firms have a stronger competitive incentive to do the same. The result of this is an industry where all firms collect real-time data on each other and on market characteristics. This transparent market would not have occurred without the incentive of gaining an ‘algorithmic competitive advantage’.

Frequency of interaction

5.27 Price adjustments, and the detection of price adjustments, require a significant amount of time and resources in brick-and-mortar retailers. With the use of algorithmic pricing, firms can reprice their products many thousands of times per day.³⁷ As described earlier, even back in 2013, Amazon has implemented millions of price changes per day, whereas Best Buy and Walmart were only able to adjust prices approximately 50,000 times during a month. As a result, when firms are tacitly colluding using algorithmic pricing, they will be able to detect and respond to deviations from collusion almost immediately.

5.28 In the extreme case if there is no delay before punishment then there is no benefit to deviation and coordination can be established regardless of the discount rate.

Calculation of optimal price

5.29 A further reason why algorithms may make tacit coordination more likely is that they may be more capable or efficient at calculating the profit-maximising tacit coordination price in the absence of an explicit agreement. As Mehra (2016) notes, there may be ‘instances in which humans would be cognitively incapable of assessing their competitors’ responses.’³⁸ In some cases, an

³⁶ OECD Secretariat (2017), [Algorithms and Collusion - Background Note by the Secretariat](#).

³⁷ Profitero (2013), [Profitero Price Intelligence: Amazon makes more than 2.5 million daily price changes](#).

³⁸ Mehra, S (2015), ‘[Anitrust and the Robo-Seller: Competition in the Time of Algorithms](#)’, *Minnesota Law Review*, Vol.100 (forthcoming).

algorithm may be better able to calculate the profit-maximising price, taking into account the predicted responses of their competitors.

Incentive compatibility and choice of pricing algorithm

- 5.30 Anticipating the discussion on pricing algorithm simulations, we know that it is possible for firms using simple pricing algorithms to reach and sustain collusive outcomes, at least in simple, highly stylised market environments. The algorithms do not even have to be highly complex. Simple ‘win-continue lose-reverse’ algorithms and ‘match low price’ (tit-for-tat) algorithms have been shown to be capable of sustaining collusion. (in fact, complex deep learning algorithms often fail to collude, especially when competing against humans as well as algorithms. However, these collusive outcomes are vulnerable to disruption, if there is uncertainty (‘noise’) in the observed prices or changes in the market (e.g. firms’ costs and demand conditions).
- 5.31 The deeper question is whether firms would find it in their interest to implement and stick to such a pricing algorithm. It is not at all clear that simply replacing a human price-setter with a pricing algorithm would solve the central problem that businesses seeking to coordinate prices face (i.e. that it is in their short-term interest to ‘cheat’ on the coordinated price, while in their longer term interest to maintain coordination), in the wider game where firms are choosing pricing algorithms (i.e. deciding whether to sacrifice short-term profits by giving up control of their prices to an algorithm that keeps prices high).
- 5.32 Salcedo (2016) explicitly models this wider game. In his model, firms simultaneously and independently commit to a pricing algorithm in the short run, and compete in a product market in which customers with unit demand arrive randomly over time. Firms must commit to a pricing policy in the short term because it takes time to revise an algorithm, but over time firms get (stochastic) opportunities to a) successfully infer or ‘decode’ others’ pricing algorithms and b) revise their own pricing algorithm.
- 5.33 Salcedo finds that, if customers arrive frequently, and revision opportunities are infrequent, then any equilibrium of his model will have long-run industry profits that will be arbitrarily close to the monopolist level.³⁹
- 5.34 Salcedo’s result is underpinned by very strong assumptions. These include perfect ex post observability, not only of all market outcomes (prices, sales, customer arrivals), but also of rivals’ algorithms, including information on how

³⁹ Salcedo, B. (2016), [Pricing Algorithms and Tacit Collusion](#).

that algorithm would respond to hypothetical outcomes that haven't yet been observed in the market. This is because, for the result to hold, firms need to be able to interpret 'proposals' to raise prices embedded within rivals' algorithms, such as a feature to match price increases from particular competitors. It is not clear that these assumptions would be satisfied in real markets, particularly where firms are using machine learning pricing algorithms that are effective "black boxes", or where there is opaque and personalised pricing. Furthermore, it could be argued that this process of decoding rivals' pricing algorithms could be understood as more akin to a form of explicit private communication between the firms (or as a public announcement of pricing intentions, if customers could also 'decode' pricing algorithms), rather than a model of tacit coordination.⁴⁰

Conclusions on likelihood of tacit coordination

- 5.35 Of the theories of harm outlined above, we consider that the hub-and-spoke model is likely to present the most immediate risk. This is because it simply requires firms to adopt the same algorithmic pricing model. Additionally, third party providers of pricing algorithm services may be a natural (and potentially 'unwitting') 'hub' for hub-and-spoke collusion.
- 5.36 The predictable agent and autonomous machine models of coordination could also occur if the pricing algorithms were sufficiently technologically advanced and in widespread use. It is unclear how likely they are to materialise at this point. However, these concerns may become more important in the future.
- 5.37 Anticipating our discussion in section 8 on risk factors for coordination, we think that algorithmic pricing is more likely to facilitate collusion in markets which are already susceptible to coordination. For these 'marginal' markets, the increasing use of data and algorithmic pricing may be the 'last piece of the puzzle' that could allow suppliers to move to a coordinated equilibrium. However, in our tentative view, it seems less likely than not that the increasing use of data and algorithms would be so impactful that they could enable sustained collusion in markets that are currently highly competitive, or those with very differentiated products, many competitors, and low barriers to entry and expansion.

⁴⁰ Schwalbe, U. (2018), [Algorithms, Machine Learning, and Collusion](#).

6. Pricing algorithm simulations and experiments

- 6.1 There is a sizable multi-disciplinary literature that studies the performance and interactions between pricing algorithms, drawing on insights from operations research, computer science, and economics.⁴¹ Most of these studies are based on simulations, rather than on applications to real-world markets. This literature is relevant to understanding the competition effects of algorithms because it can help demonstrate the conditions under which coordinated outcomes may emerge.
- 6.2 To get a broad overview of this literature, studies can be categorised by the type of pricing problems that they examine. There are many possible combinations of factors and assumptions, each of which would set up a different problem or model. Many of these pricing problems are too difficult to study analytically (i.e. by obtaining exact solutions), but instead have been subjected to numerical approaches using pricing algorithms and simulations. The main factors and assumptions determining the type of pricing problem are:
- (a) Nature and knowledge of demand – whether demand is i) stable over time and is known by sellers (which is the assumption in classic oligopoly models studied by economists), ii) stable over time but unknown to sellers, or iii) varies over time and unknown to sellers.
 - (b) Finite or infinite selling period – whether there is a finite inventory or selling period (and if so, whether there is any replenishment of inventory after each period), as opposed to an infinite inventory and selling period.⁴²
 - (c) Competition – whether there is a single seller or multiple sellers and, in the context of multiple sellers, whether only one, a few, or all the sellers are using algorithms⁴³ as opposed to some simpler method of pricing.

⁴¹ For a relatively up-to-date review of this literature, see den Boer (2015), '[Dynamic pricing and learning: Historical origins, current research, and new directions](#)', *Surveys in Operations Research and Management Science Vol.20*, pp1-18.

⁴² The latter assumption is more common in oligopoly models studied by economists. However, the literature studying algorithms for finite inventory and selling periods seems to be far more developed, because it reflects the pricing problem faced by airlines and hotels when deciding how to price a (short-term) fixed supply of services that are perishable. Airlines and hotels were two of the earliest sector to develop dynamic pricing algorithms that reflect the marginal value of the remaining inventory (setting higher prices when it seems like demand will exceed remaining supply) and the time left to sell them (all other things being equal, including customer willingness to pay, setting lower prices when there is less time remaining).

⁴³ Adding multiple learners can greatly complicate the problem, even without considering the issues of strategic interdependence. For example, Q-learning (a prominent type of reinforcement learning) is known to converge to the optimal policy when the function to be learned is a stationary process (i.e. it does not shift over time). But

Studies based on monopolies are usually focused on the performance of pricing algorithms in helping sellers to discover optimal prices when demand is unknown or when demand can vary over time (or both), either by using data on customers or by experimenting with prices⁴⁴, in finite inventory or selling time contexts.

As for studies of pricing algorithms with multiple sellers, taking a broad view, there is a large literature that is relevant to the performance of pricing algorithms in oligopolistic markets. In many models of oligopoly firms, the incentives faced by firms in deciding whether to coordinate or to compete is analogous to a prisoner's dilemma.⁴⁵ Iterated prisoner's dilemma is a widely-studied game in economics and biology, and the ability of various static and evolutionary algorithms to reach and sustain cooperation has been studied extensively.⁴⁶

- (d) Strategic consumers – whether customers' willingness to pay is impacted by the sellers' past actions (such as when selling to repeat customers with reference price effects) or sellers' expected future actions (such as when selling to customers who can strategically anticipate sellers' lower prices in future, as opposed to myopic customers that make a single 'buy-or-leave' decision when they arrive).⁴⁷ Various assumptions about consumer behaviour can also be introduced to increase realism, for instance by assuming that only a portion of customers check all offers on the markets and select the best one in each period, whilst other 'captive' customers may only check a handful of offers infrequently.

when two Q-learners are pitted against one another, each creates a non-stationary environment for the other. There are no theoretical guarantees of convergence in this case. See Kephart et al. (2000), '[Dynamic pricing by software agents](#)', *Computer Networks Vol.32*, pp731-752.

⁴⁴ When sellers are experimenting with prices to try and learn about demand, they are solving a 'multi-armed bandit problem' in which they have to trade off 'exploration' (using some of the limited inventory or time period to learn about demand) and 'exploitation' (setting the optimal price and maximising profit based on whatever limited knowledge of demand they have in the remaining time available).

⁴⁵ Prisoner's dilemma (PD) is a game in which players choose simultaneously between two actions: cooperate or defect. In PD, defecting is a dominant strategy for every player (i.e. each individual's payoff from defecting is higher than the individual payoff from cooperating, regardless of the other players' actions), but the joint payoff from mutual cooperation is higher than the joint payoff from mutual defections.

⁴⁶ Arguably, the performance and abilities of algorithms to sustain coordination has been studied at least since Axelrod's iterated prisoner's dilemma tournaments in the early 1980s. See Axelrod, R. (1984), *The Evolution of Cooperation*.

⁴⁷ This affects whether, for instance, sellers could gain from committing to a preannounced pricing policy for the entire selling season, without considering the information that it will gain from realised future sales. In a setting where customers are myopic (and there is no competition), sellers would never benefit from tying their hands in this way.

(e) Multiple products and product differentiation – whether sellers are selling a single, homogenous product, or whether they are selling multiple products (and, if the latter, whether these are substitutes, complements, or independent, and whether they are horizontally or vertically differentiated). Models with multiple products are needed to study, for example, i) joint pricing-inventory or pricing-production problems, in which sellers can replenish their inventory at certain points in the selling season but must decide what to stock and how to price them, and ii) cross-selling, upgrading and upselling, in which customers make an ‘initial’ purchase decision but sellers can use the information revealed to influence the customers’ ‘final’ purchasing decision.⁴⁸

6.3 In addition to the range of pricing problems, the literature has also tested a range of algorithms. Giving a full summary of the range of results about the interactions of different pricing algorithms simulated in the literature is beyond the scope of this paper. Instead, we present some high-level observations from our review of this literature:

- (a) Simple reinforcement learning algorithms can achieve coordinated (i.e. cooperative) outcomes under conditions of perfect information, but it is much harder to do so if there is even a small amount of noise or uncertainty.⁴⁹ Similarly, coordinated outcomes can also be disrupted by small changes in the market, such as firms’ costs and demand.⁵⁰
- (b) There is a strand of literature that uses Q-learning agents for prisoner’s dilemma and other games, which is aimed at devising algorithms to improve cooperation.⁵¹ This literature highlights that whilst algorithmic

⁴⁸ See Chen and Chen (2014), ‘[Recent Developments in Dynamic Pricing Research: Multiple Products, Competition, and Limited Demand Information](#)’, *Production and Operations Management* 24(5), pp704-731

⁴⁹ For instance, see Miller, J.H. (1996), ‘[The Coevolution of Automata in the Repeated Prisoner’s Dilemma](#)’, *Journal of Economic Behaviour & Organisation*. Miller (1996) experimented with a genetic algorithm in repeated prisoner’s dilemma under three information conditions (perfect information, 1% noise, and 5% noise, where noise is the probability that an action is misreported), and found that average payoffs for the evolved automata under the noisy conditions plateaued at a lower level than under perfect information.

⁵⁰ Izquierdo, S.S., and Izquierdo, L.R. (2015), ‘[The “Win-Continue, Lose-Reverse” rule in Cournot oligopolies: robustness of collusive outcomes](#)’, *Advances in Artificial Economics*, pp33-44. Izquierdo and Izquierdo (2015) show that a simple type of reinforcement learning (Win-Continue, Lose-Reverse) leads to a cooperative outcome in a Cournot oligopoly game, but that this is not robust to small independent perturbations in the firms’ cost or demand functions. In markets with such perturbations, WCLR converges instead to a non-coordinated Cournot-Nash equilibrium.

⁵¹ For instance, see de Cote, E.M., Lazaric, A., and Restelli, M. (2006), ‘[Learning to cooperate in multi-agent social dilemmas](#)’, *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pp.783-785. See also Crandall, J.W. and Goodrich, M.A. (2011), ‘[Learning to Compete, Coordinate, and Cooperate in Repeated Games Using Reinforcement Learning](#)’, *Machine Learning* 82(3) pp.281-314.

coordination is possible, it is not necessarily straightforward as many reinforcement learning algorithms converge instead on the competitive equilibrium. It often appears that results are often very specific either to the set-up of the problem studied or the exact formulation of the pricing algorithms tested. In addition, often cooperation can be achieved with simple algorithms, while more complex algorithms based on deep learning often fail to cooperate.⁵²

- (c) Many studies of algorithmic pricing involve self-play, where all the players have the same algorithm, but in real-world markets where rivalry has not been fully suppressed by explicit coordination, we may expect competitors to experiment with different pricing algorithms and approaches. If players use different pricing algorithms, each of which could be learning over time, this greatly increases the complexity and difficulty of establishing coordination.⁵³ This could indicate that we should be more concerned about tacit coordination if algorithms are very prevalent, and that if more firms utilise the same pricing algorithm in the same market, it makes it more likely that the market will move to an outcome where prices are higher.
- (d) Furthermore, although these studies show that coordination can be achieved, often this is by using simple algorithms that result in behaviour which can be exploited by competitors. In effect, these studies effectively assume that the firms would all implement and stick to this algorithmic strategy, which assumes away the crucial question of whether firms would have the incentive to do so in practice. It is relatively simple to choose model assumptions and algorithmic strategies that will give rise to a tacitly coordinated outcome in simulations, if we ignore game theoretic concerns about whether players have an incentive to follow these strategies (i.e. incentive compatibility).

⁵² Schwalbe, U. (2018), 'Algorithms, Machine Learning, and Collusion'.

⁵³ For example, see David Foster, "Algorithms and Price Collusion. Or How I Learned to Stop Worrying and Love AI" forthcoming in Competition Law Insight, <https://www.competitionlawinsight.com/> and <http://www.frontier-economics.com/publication/algorithms-price-collusion/>. Foster used a differentiated Bertrand model with several firms using a 'trial and error' pricing algorithm. If all the firms in a differentiated Bertrand market used such an algorithm, Foster found that the firms would reach a tacitly collusive outcome. However, if only one large firm is using this algorithm, coordination fails to be established.

7. Algorithms and personalised pricing

- 7.1 Concerns about increasing availability of data and use of pricing algorithms are not limited to their potential to exacerbate collusion. A second set of concerns is that, in combination with the growth of ‘Big Data’, they might lead to personalised pricing.
- 7.2 We define personalised pricing as the practice where businesses may use information that is observed, volunteered, inferred, or collected about individuals’ conduct or characteristics, to set different prices to different consumers (whether on an individual or group basis), based on what the business thinks they are willing to pay.
- 7.3 In many cases, personalised pricing can be beneficial – for example the ability to offer targeted discounts might help new entrants to compete particularly in markets with switching costs, and could be output-expanding. On the other hand, there may be situations where personalised pricing can lead to consumer harm.
- 7.4 The conditions under which competition authorities might be more concerned about personalised pricing were outlined in an OFT economics paper in 2013⁵⁴, and we do not repeat them here. Instead, this section focuses particularly on the ways in which firms may be using ‘Big Data’ and algorithms to facilitate personalised prices, and whether this could undermine collusive outcomes.
- 7.5 We note that it is already possible for businesses to price discriminate, using a few observable customer characteristics (e.g. student discounts, or third-degree price discrimination) or by providing options which induce customers to self-select different effective prices (e.g. quantity discounts, or second-degree price discrimination). However, the increasing availability of data and use of sophisticated pricing algorithms, particularly by online retailers, raises the possibility that such retailers would be able to engage in highly personalised pricing, effectively sorting customers into ever finer categories. In the extreme, the outcomes of highly personalised pricing may approach those of perfect or ‘first-degree’ price discrimination, in which every customer is offered an individual price equal to their maximum willingness to pay.

⁵⁴ OFT (2013), [The Economics of Online Personalised Pricing](#).

7.6 We first summarise some of the evidence about the existence of personalised pricing, before discussing the interaction and tension between tacit coordination and personalised pricing.

Evidence of personalised pricing in practice

7.7 There is evidence that the majority of consumers dislike online personalised pricing. For example, there was a backlash in 2000 when Amazon.com varied the prices of their DVDs, allegedly based on previous browsing patterns.⁵⁵ This practice was found after customers discovered that they could buy products at a lower price if they stripped their computer of the electronic tags that identifies them as a regular customer. Following this, Amazon denied personalising prices, and stated that the price deviations were totally random discounts to test how sales would change in response to price changes. They refunded all customers who received higher prices.⁵⁶

OFT Personalised Pricing Call for Information in 2012

7.8 In 2012, the Office of Fair Trading (OFT) launched a call for information to improve its understanding of how the use of consumers' information is affecting online markets. The businesses they discussed this with stated that they had no desire to identify individual consumers, and were aware of the potential adverse consumer reaction to actual or perceived invasions of their customers' privacy.

7.9 The OFT also conducted its own research, which found no evidence of prices being set on the basis of individual consumer profiles by Amazon or any other company, as opposed to a broader group, or type, of consumer.⁵⁷

7.10 The OFT noted that consumers and media commentators find it difficult to distinguish between personalised pricing and other forms of price discrimination. Online prices can vary rapidly and consumers may think they are being offered a price based on information collected about them personally.

CMA research on personalised pricing in 2017

7.11 To investigate the potential for algorithms to facilitate personalised pricing, we expanded on the OFT 2012 experiment. We compared the prices of a

⁵⁵ See BBC News article from 2000, [‘Amazon’s old customers “pay more”](#)’.

⁵⁶ See Amazon’s press release on this, [Amazon.com Issues Statement Regarding Random Price Testing](#).

⁵⁷ OFT (2013), [Personalised Pricing: Increasing Transparency to Improve Trust](#).

selection of varied products, across different operating systems, when accessed directly or through an affiliate website (e.g. a digital comparison tool or cashback site), and against logged in customer profiles, to determine if there was any evidence of personalised pricing.

- 7.12 There was no evidence of pricing being different or personalised for different consumers. There were examples of different consumers being shown different search results on retail websites, including different numbers of results or a different order of results. For more detail, refer to Annex 1.

European Commission (DG Justice) Consumer market study on online market segmentation through personalised pricing/offers in the European Union (2018)

- 7.13 DG Justice commissioned research on the extent of personalised pricings and office in EU member states between December 2017 and November 2017.⁵⁸ This research included:

- (a) a mystery shopping exercise in four online markets (airline tickets, hotels, sports shoes, and TVs);
- (b) an online behavioural experiment designed to assess consumers' ability to recognise and respond to online personalisation, as well as the effect of disclosure messages about personalisation on consumer awareness and responses; and
- (c) a survey of consumer awareness and attitudes to online personalisation practices.

- 7.14 Focusing on the mystery shopping exercise, the researchers found evidence of online personalised ranking of offers.⁵⁹ 61% of the 160 e-commerce websites visited were found to do this, either based on information about the shoppers' access route to the website or based on past browsing behaviour. However, they did not find evidence of systematic and consistent personalised pricing, across the eight EU member states and four markets.

⁵⁸ European Commission (2018), '[Consumer market study on online market segmentation through personalised pricing/offers in the European Union](#)'.

⁵⁹ This practice of presenting different results or the same results in a different order to different consumers that have made the same search, on the basis of information about the consumer's characteristics is also known as 'search discrimination' or 'price steering'.

Price differences were only observed in 6% of tests, and the median price difference observed was less than 1.6%.

Hannak et al (2014)

- 7.15 Hannak et al. (2014)⁶⁰ surveyed 16 popular e-commerce websites (10 general, 6 hotel and car rental) to measure price discrimination and price steering in the US. They used two methodologies.
- 7.16 First, they recruited real users living in the US from Amazon's Mechanical Turk service, and asked them to complete a search for prices after configuring their browser so that their search was routed through the authors' HTTP proxy, which could make simultaneous and identical 'control' searches without any of the users' cookies. This method controls for differences by geolocation (all searches come from the same IP address of the authors' proxy server) and by time.
- 7.17 Using this methodology, the authors found that:
- (a) For some websites, users were consistently receiving results in a different order relative to the control searches. For example, they found that Sears appeared to be ordering search results for users so that *cheaper* products are displayed near the top, compared with the control searches.
 - (b) A small set of sites (Home Depot, Sears, and many of the travel sites) displayed different prices to users relative to controls, for a small but significant proportion of products tested (between 0.5 to 3.6% of products), and the average difference in prices in these cases were hundreds of US dollars higher for users relative to control searches.⁶¹ In addition, the authors noted that some users appeared to experience personalisation across multiple websites.
- 7.18 Although the first methodology demonstrated that websites personalise results for different users, it was not possible using this method to determine why and on what basis the personalisation was set, as real-world user profiles had too many potentially confounding variables. Therefore, the authors also conducted controlled experiments by creating false accounts for which some

⁶⁰ Hannak et al. (2014), '[Measuring Price Discrimination and Steering on E-commerce Web Sites](#)', *Proceedings of the 2014 conference on internet measurement conference*, pp305-318.

⁶¹ The authors did not report the size of the price difference in percentage terms, but in Figure 4 of Hannak et al. (2014), the authors show an example for a hotel in Paris, which cost \$633 for one user compared with \$565 for the control search, an increase of 12%.

variables were changed but were otherwise identical. The variable features tested were: browser, OS, account log-in, click history, and purchase history.

7.19 They found that:

- (a) Only one retail site (Home Depot) and none of the rental car sites revealed personalised prices based on the user features tested.
- (b) Cheaptickets and Orbitz offered logged-in members reduced prices for hotels.
- (c) Expedia and Hotels.com conducted randomised A/B testing on users, including steering one group of consumers towards more expensive hotels.
- (d) Priceline altered hotel search results based on the user's history of clicks and purchases. Users that clicked on or reserved low-price hotel rooms receive slightly different results in a much different order, compared to users who click on nothing, or click/reserve expensive hotel rooms.
- (e) Travelocity offered lower prices for hotel rooms to customers using iOS devices. Also, users browsing with Safari on iOS (a mobile device) received slightly different hotels, in a much different order, compared with users browsing with Chrome on Android, Safari on OS X, or other desktop browsers.
- (f) Home Depot also personalised results for mobile users. On most days, there was almost zero overlap between the results displayed to desktop and mobile browsers, but on a few days the results are identical for all browsers. The pool of results served to mobile browsers contained more expensive products overall.

Mikians et al. (2012)

7.20 Mikians et al. (2012)⁶² also conducted controlled experiments, and tested the effect of varying: browser, OS, geolocation (using proxy services in east and west coast US, Germany, Spain, Korea, and Brazil), browsing history (using false personas to build up a browsing history that mimics the behaviour of affluent or budget conscious people), and origin URL (i.e. which website you

⁶² Mikians et al. (2012), '[Detecting price and search discrimination on the Internet](#)', *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*.

were on before visiting a vendor's website). They looked at 600 products, across 35 product categories from 200 distinct vendors.

7.21 The authors' methodology for building up a browsing history is particularly instructive. Mikians et al. (2012) used Audience Science to research generic traits and browsing habits for budget conscious and affluent consumers. For instance, budget conscious consumers visit price aggregation and discount sites more often than average, and affluent consumers visit high-end luxury products, automotive resources, and community personals sites more often than average. Using Alexa and Google's lists of most popular sites, the authors selected appropriate sites for each profile, and built up the profile by visiting these sites for a week. During this training session, the authors permitted tracking and disabled all blocking, so that the profiles may be tracked by third party aggregators and ad networks (such as Google Analytics/Ad Services and DoubleClick) that have a presence on many sites, and can combine information about these visits to build up a profile of the user.

7.22 Mikians et al. (2012) found:

- (a) No evidence of price or search discrimination for different OS and browsers.
- (b) Price differences based on geographic location of customer, primarily for digital products (such as ebooks and video games) of up to 166%. They also observed price differences for Staples's website (an office products seller) when queries are sent from different locations within Massachusetts, USA. However, these differences could be due to digital rights costs and competition rather than price discrimination.
- (c) Evidence of search discrimination (but not price discrimination) between affluent, budget conscious, and 'clean' profiles (with no browsing history) on online hotels and ticket vendor websites. The price of the different products shown to affluent personas are up to four times higher than for budget conscious personas.
- (d) They also found price discrimination depending on the channel (or origin URLs). For some product categories, when a user visits a vendor site via a discount aggregator site, prices can be up to 23% lower than what is available when visiting the vendor's site directly.

Austrian Chamber of Labour (Arbeiterkammer Wien) studies in 2016 and 2017

- 7.23 For every day in a week in 2016 (between 2pm and 3pm on Monday to Friday, and between 9am and 10am on Saturday and Sunday), AK Wien looked up 36 prices for various products on more than 28 devices (including desktops, laptops, smartphones, and iPads) in the 9 federal capitals of Austria over a period of 5 days. They also did this in Dusseldorf (in Germany). The products examined included furniture, flights, shoes, and hotels. They tested for changes in price by type of device, geographical location, and time.
- 7.24 They found no differences in price by type of device. For some products, they found that prices varied over the course of the week, and that there were differences in prices on online store websites accessed in Austria and in Germany.
- 7.25 In 2017, they repeated this study, examining 33 prices on different online shops (Amazon, Lufthansa, AirBerlin, Austrian Airline, Opodo, booking.com and Heine) on 20 different devices (stationary desktops, laptops, notebooks, iPads, smartphones and iPhones) every Tuesday, Thursday and Saturday, from the 14th to the 25th of March 2017. The devices were located in various places in Austria and one device (laptop) was located in Dusseldorf, Germany. Between 3 and 9 products were checked on each site.
- 7.26 This time, they found some differences in prices between devices, mainly for travel,⁶³ but it was not the case that prices were always higher for one type of device. They also noted that, for Opodo, some products were offered on some devices but not on other devices, and this changed daily. Again, they also found differences in price across time and between Austria and Germany.

What is the interaction between collusion and personalised pricing?

- 7.27 Most of the theories of harm considered in section 5 assumed that algorithmic collusion results in the same price being offered to every consumer, such that consumers have no real choice among competitors (i.e. that prices are not personalised). However, as discussed above, algorithms could make personalised pricing less resource intensive and more accurate.

⁶³ The range of price differences for the product with the largest price difference: Air Berlin €5 - €10; Austrian Airlines €30 - €80; Opodo €28 - €167; Booking.com up to €154.35 (for hotels in Madrid) and up to €66.35 for hotels in Hamburg, but no differences for other cities. For Lufthansa and Heine, there were no price differences depending on device.

7.28 Therefore, a key question is whether collusion, either explicit or tacit, could work in tandem with personalised pricing. In this section, we discuss the compatibility of the ‘traditional’ factors necessary for coordination and for personalised pricing, and present our views of how personalised pricing and tacit coordination could exist in tandem using algorithms.

Conditions for coordination and personalised pricing

7.29 For firms to be able to coordinate, they should be able to:

- (a) reach a common understanding of and implement coordinated terms (pricing structures);
- (b) find and agree on a coordinated outcome and share of profits which is acceptable to all participants and preferable to what participants would receive under competition (incentive compatibility and allocation structures);
- (c) detect and credibly punish deviations (enforcement structures); and
- (d) prevent outsiders from disrupting the coordination (e.g. barriers to entry; overcoming buyer resistance).

7.30 For firms to be able to charge personalised prices, they should be able to:

- (a) observe or possess information about each customer’s willingness to pay; and
- (b) prevent resale between customers or customer segments.⁶⁴

Explicit collusion and personalised pricing are compatible but unlikely to occur together

7.31 In principle, it is possible for firms to explicitly collude and engage in personalised pricing. The necessary conditions for successful collusion and personalised pricing are compatible, but probably unlikely in practice.

7.32 If firms successfully establish sustainable and explicit collusion, then they can exploit their joint monopolist position by using personalised pricing. Once

⁶⁴ For perfect (first-degree) price discrimination, in which firms charge each customer their maximum willingness to pay, firms would need complete information about customer’s WTP. Also, there are two further necessary conditions: 1) no competition (otherwise the competitors will undercut the firm charging maximum WTP; and 2) a full set of pricing instruments (so that firms can set marginal price equal to its marginal cost, and extract each customer’s entire surplus with a fixed fee).

competition has been suppressed then, provided that the other necessary conditions for personalised pricing are met (e.g. no resale, sufficient information about customers), the cartelists can share and use data about each customer's willingness to pay and pricing algorithms to set personalised prices in order to extract the maximum consumer surplus (but also to detect, respond to, and potentially deter deviations and/or entry). In effect, the firms act as a joint monopolist implementing first degree price discrimination.

- 7.33 Clearly, it is more difficult and complex to sustainably coordinate across many personalised prices and to enforce the terms of the coordination. However, it seems at least conceptually possible to overcome these difficulties if firms were supported by sufficiently sophisticated data and algorithms, as well as explicit communication and sharing of information. Whether firms could actually do so is an empirical question about the quality and availability of data and algorithms.
- 7.34 Nonetheless, we expect that there are likely to be relatively few retail markets in which there could be both explicit coordination and personalised pricing. Regardless of whether firms are using pricing algorithms, for both collusion and personalised pricing to coexist, all the 'traditional' conditions for both perfect price discrimination and collusion should be satisfied, and this is quite unlikely.⁶⁵ In addition, we suspect that, particularly in retail markets, there may be a tension between a) the transparency and level of information needed to explicitly coordinate over many personalised prices, and b) the opacity needed to evade detection by competition authorities and to prevent customer resistance, particularly to personalised prices.⁶⁶ There would need to be a very large asymmetry between cartelists and customers/regulators in technical ability and access to information about prices and transactions.

Tacit coordination and personalised pricing are very unlikely to occur together

- 7.35 Without explicit communication and sharing of information, if there are many differentiated products and personalised prices, then it appears far more difficult to reach a common understanding of the terms of coordination.
- 7.36 It is likely to be very difficult, even if there is sufficient objective and observable information about customers' willingness to pay that would allow each firm to independently derive common focal prices for each customer

⁶⁵ Prices that vary with each customer are less likely to be transparent to competitors.

⁶⁶ Certainly this would seem to require significant intention, cooperation, and effort by the coordinating firms.

using pricing algorithms, simply due to the number of prices leading to increased opportunities for errors and miscoordination.

- 7.37 Also, in retail markets with personalised pricing, there is often limited transparency or public information about transactions or the actual prices paid by customers. Without explicit sharing of information, it seems that firms would not be able to detect and credibly punish deviations, and firms could make secret offers to customers that undercut any tacit price level.
- 7.38 For these reasons, any tacitly collusive outcome also is also likely to be highly unstable, and it appears much less likely for firms to reach a tacit coordinated outcome whilst setting personalised prices.

Ezrachi and Stucke model of tacit coordination and personalised pricing

- 7.39 Ezrachi and Stucke (2017)⁶⁷ propose a model under which firms are able to tacitly collude and apply personalised pricing to different customer groups. In this model, firms offer a completely transparent price to all customers, which can be set supra-competitively. They then identify the high-value customers, either using an algorithm or other customer features (e.g. by identifying those coming directly to the website instead of through a search engine). Ezrachi and Stucke suggest that the firm could offer these high-value customers personalised prices through a “secret”, completely un-transparent, deal. These could be personalised discounts (relative to a high initial price), or alternatively higher prices through “drip pricing”. In particular, targeting customers with higher personalised prices enables the firms to capture additional value from high value customers.
- 7.40 Ezrachi and Stucke describe this outcome as providing customers with the ‘worst of both worlds’. This is because once a customer has been placed into either the loyal or low-value groups, they are offered prices above the competitive rate (via collusion). Customers who are high-value have their consumer surplus extracted via personalised pricing. Both loyal and low value customers pay a price that is publicly advertised (thus firms can collude on this price). This, however, assumes that firms are able to consistently group customers into these three groups,⁶⁸ and that customers do not change groups.

⁶⁷ Ezrachi, A, and Stucke, ME (2016), *Virtual Competition: The Promise and Perils of the Algorithm-Driven Economy*.

⁶⁸ It is not clear why Ezrachi and Stucke propose these three customer groups, nor how value (which we assume refers to the amount that customers are willing to spend) and loyalty (which we assume refers to customers' preference for one casino over another) interact.

- 7.41 In our view, the main drawback to this model is that it is unclear how collusion can be sustained if firms are able to provide secret offers to some customers, for example, through direct email offers. As discussed previously, monitoring is vital to the stability of a collusive agreement.
- 7.42 Furthermore, there does not seem to be any reason in the model why secret personalised offers could not also be extended to the low-value and loyal groups of customers.⁶⁹ If firms can cheat on a coordinated outcome without being detected (e.g. by offering a lower personalised price that other firms do not know is being offered), then this will always be profitable, no matter what category of customer. Even supposing, as Ezrachi and Stucke do, that collusion could be sustained for the low-value and loyal groups, and the ability to make secret offers is limited to high-value customers if most industry profits are from the high-value customers, then the ability to offer secret deals will restore competition for a substantial part of the market.

Conclusion on the interaction of coordination and personalised pricing

- 7.43 The increasing availability of data, and the sophisticated use of pricing algorithms, increases the scope for tacit coordination or personalised pricing. However, in our view, it is unlikely for both tacit coordination and personalised pricing to occur within the same market. The ‘traditional’ conditions that facilitate tacit coordination make it harder to engage in highly personalised pricing because price comparisons are easy for customers, and the increasing use of data and algorithms does not change this.
- 7.44 As to the question of which is more likely to occur in any given market, tacit coordination or personalised pricing, this will depend on – among other things - the extent to which the necessary conditions outlined in paragraphs 7.31 and 7.32 above are fulfilled. We note that, in the abstract, personalised pricing appears to have more difficult information and computational requirements than tacit coordination (which, as we discussed in previous sections, can be implemented very easily if we set aside the question of incentive compatibility). It will also depend on the expected cost, including implementation costs but also the perceived likelihood and consequences of detection, relative to increased profit from collusion or personalised pricing.

⁶⁹ A high collusive price may continue to apply for the low-value and loyal groups, if there is at most one sophisticated firm in the market that can implement personalised pricing for all groups (i.e. as the result of technological asymmetries and capacity constraints, due to the costs of implementing personalised pricing beyond a small group of high-value customers).

8. Features which might raise competition concerns

- 8.1 This section draws together some tentative conclusions on what features of the market or of the algorithms themselves might lead to greater concern about algorithmic pricing.
- 8.2 We first outline risk factors which might raise concerns about algorithmic pricing leading to coordinated outcomes. We then discuss risk factors relating to personalised pricing.

Risk factors for coordination

- 8.3 Below, we list some traditional risk factors for price co-ordination, with details regarding how algorithmic pricing and online markets in general could result in these creating more risk of harm to consumers.
- (a) **Concentrated markets.** Because of algorithmic pricing, both explicit and tacit co-ordination could occur in less concentrated markets. This is because algorithms can collect information about more competitors at a faster rate than humans. Therefore, deviations can be detected from more firms and punishment strategies could be implemented more rapidly.
 - (b) **Market Transparency.** Increasing availability of data and use of pricing algorithms can increase market transparency (especially online), even if there are many products with complex offers. This is because algorithms can scrape data from many websites more quickly than humans would be able to. Therefore, price deviations can be detected more quickly. Not only do algorithms lead to greater information on competitors' actions and customers but they can lead to simpler and more predictable pricing behaviours.⁷⁰
 - (c) **Frequency of interaction and price setting.** Pricing algorithms allow firms to set their prices automatically. This means that whenever a price change occurs, competitors can respond by undercutting or matching very quickly with low or zero menu cost. Therefore, the short-term return for a firm that lowers its price below the market price may be very small, which would discourage price wars. Further to this, an algorithm could allow firms to test their competitors' responses to changes in price during periods of low demand, because price feedback is so quick and multiple rounds of price change can occur very quickly even when there are no

⁷⁰ Algorithms can gather competitor price information, gather other online data on website use, and cause prices to more quickly and obviously track a market signal.

customers present in the market. Although the firm that has been undercut in this price war would not face a real punishment, because little to no sales have been made, the algorithm would learn about the likely effects of future price competition and may be discouraged from engaging in price wars at busier periods.

(d) **Low buyer power or small regular purchases.** Online markets operate 24 hours a day, and sometimes internationally, therefore there can be a high frequency of purchase with low buyer power. However, as described previously, there may be opportunities for customers to use algorithms to form buying groups, increasing their buyer power.

8.4 Algorithms could potentially increase the chance that tacit coordination occurs in ways that go beyond traditional risk factors:

(a) An algorithm could monitor prices, introduce parallel conduct (e.g. follow the price leader), signal to competitors about intentions or just learn to coordinate.

(b) An algorithm could increase the stability of a cartel by increasing barriers to entry, if it is able to identify and quickly target customers who are most likely to buy from a new entrant (a form of personalised pricing).

(c) Firms using the same algorithm or the same data set (which means the algorithm learns/adapts in the same way) may act in parallel.

8.5 The mechanisms by which algorithms could have an additional impact beyond 'traditional' risk factors are quite speculative and are likely to be difficult to evidence. Instead, the main impact of increasing use of data and algorithms appear to be that it can exacerbate 'traditional' risk factors, such as transparency and the speed of price setting.

8.6 Algorithmic pricing is more likely to facilitate collusion in markets which are already susceptible to (human) coordination. For these "marginal" markets, the increasing use of data and algorithmic pricing may be the 'last piece of the puzzle' that could allow suppliers to move to a coordinated equilibrium.

8.7 Factors which could give competition authorities an indication of whether a price-setting algorithm may result in tacit coordination include:

(a) **The time horizon of a reinforcement learning algorithm's objective function.** It is plausible that if the objective function is very short-term or places a large weight on short-term profits (e.g. maximise profit just on one sale) then a reinforcement learning algorithm is less likely to engage in stable coordination. For stable tacit coordination to take place, the firm

must be willing to sacrifice short term profits in favour of a longer-term, more profitable outcome. However, we note that this is may not be a reliably useful indicator of potential harm, as even simple Win-Continue Lose-Reverse algorithms, which aim myopically to maximise revenue in just the current period, could give rise to a collusive outcome in favourable conditions (e.g. other competitors using similar strategies, underlying market conditions like costs and demand are stable, etc.).

- (b) **Whether all/many competitors are using the same algorithm/objective function.** In the case of markets where intermediaries provide algorithmic price services to several competitors (a 'hub and spoke' scenario), this is closely related to how much of the market the intermediaries cover.
- (c) **What data the algorithm is using**, and in particular whether the algorithm makes use of information or data from multiple competitors, which may be a particular risk in markets where intermediaries receive data from multiple clients that are competitors.

Risk factors for personalised pricing

8.8 Personalised pricing can be beneficial in many situations (as discussed in Section 7). The OFT (2013) personalised pricing study discussed situations in which we might have greater concerns. We have considered the extent to which algorithms change these factors. The OFT study found that personalised pricing was more likely to be harmful to consumer welfare when:

- (a) There is a lack of competition in the market (i.e. a monopolist).
 - (i) Although we do not address these issues in this paper, there is a potential interaction with questions on whether, in some online markets, the collection and use of user data can lead to strong network effects which incumbents can exploit to build and maintain dominant positions. These are likely to be markets where highly personalised pricing is more feasible.
- (b) Discrimination is particularly complex (several different consumer groups) or opaque to consumers.
 - (i) This is the case for online markets, as firms can collect more data on customers in a shorter period. Additionally, using algorithms allows for this data to be processed more quickly, and allows for more data to be analysed. Firms are able to collect a greater variety of data on consumers. Therefore, they would be able to split customers into

more groups, with greater information about the customer's willingness to pay.

- (ii) Additionally, if firms can send “secret deals” to customers, for example by directly offering discounts via email, the price discrimination becomes entirely opaque. Further, given there is a wide range of potentially relevant factors which firms may use to discriminate and the potential for sophisticated firms to use combinations or the presence of multiple factors to personalise prices, it is unlikely that experiments focusing on varying one factor at a time would uncover the underlying factors for any observed personalised pricing (or indeed, any personalisation at all, if firms are cautious enough to detect and avoid attempts to uncover personalised pricing based on simplistic or single factors).
- (c) It is very costly to firms, because if firms incur significant costs to price discriminate, then they will need to recover these costs through higher prices.
 - (i) This is less likely to be the case for online price discrimination, as data is typically readily available compared to offline price discrimination. The cost associated with personalised pricing will depend on the complexity of the algorithm used.
- (d) Consumers lose trust in the market and, as a result of their lack of confidence that they are receiving a good or fair price, they may withdraw their demand or decline to participate in the market (and potentially, in other similar online markets).
 - (i) This is a big issue which we have not discussed in this paper. This is likely to be an issue with online personalised pricing, and consumer mistrust may be exacerbated by suspicions of collusion facilitated by algorithmic pricing. The extent to which this occurs will partly depend on whether customers' perceptions of price fairness⁷¹ and social norms.⁷²

⁷¹ There is a sizable multi-disciplinary literature focused on this topic. For a seminal paper in economics, see for instance Kahneman, Knetsch, and Thaler (1986a), ‘[Fairness as a Constraint on Profit Seeking: Entitlements in the Market](#)’, *American Economic Review* 76(4), pp728-741. Similarly, in the marketing literature, see for instance Xia, Monroe, and Cox (2004), ‘[The Price Is Unfair! A Conceptual Framework of Price Fairness Perceptions](#)’, *Journal of Marketing* 68(4), pp1-15.

⁷² Social norms may explain why, for instance, airline passengers do not appear to mind paying different prices from others in nearly identical seats, while Amazon faced a strong consumer response for attempting to price

discriminate in 2000 with DVDs. See Garbarino and Maxwell (2010), '[Consumer response to norm-breaking pricing events in e-commerce](#)', *Journal of Business Research* 63(9-10), pp1066-1072.

9. Further Work

9.1 We consider that there could be value in further economic research exploring the topics discussed in this paper. Some specific areas for further research could include:

- (a) **Auditing algorithms** – deep learning algorithms are often described as “black boxes”, but there is a small and emerging research community dedicated to auditing algorithms, mainly in the context of detecting discrimination based on protected characteristics (like race, sex, etc.). It may be useful to take insights from this field and apply them to auditing pricing algorithms. From an enforcement and regulatory perspective, it would be beneficial to understand further whether and if a firm could know that its algorithm is implementing a collusive outcome. For instance, if a firm observes that its profits have risen since it implemented algorithmic pricing, would it be able to determine whether this is because the algorithm has attracted new customers, increased sales to existing customers, raised prices to loyal customers, or engaged in collusion? There may be little difference between collusion and raising prices on some products that are considered less elastic for other reasons (such as loyal customers).
- (b) **Algorithmic decision rules that should be presumed to be anti-competitive** – in the case of simpler algorithms, are there certain kinds of decision rules which have no plausible rationale other than to facilitate an anti-competitive outcome? For instance, one might argue that there are no pro-competitive reasons to have decision rule in an algorithm that raises price in response to a competitor’s price increase, or a decision rule that never undercuts a competitor’s price. However, for both of these examples, these decision rules are consistent with competitive firms trying to maximise profit, and it may be too interventionist and damage the competitive process to restrict firms’ ability to set its own prices.
- (c) **Secret offers and masking** – consumers may not be helpless in response to collusion and personalised pricing, and they may be able to use countermeasures.

In the case of collusion, to what extent can customers request (potentially using an algorithm or ‘shopbot’ in an online context) secret offers from suppliers in order to undermine collusion? To what extent could customers build up and exercise buyer power through joint purchasing? From the perspective of implementing a cartel, are competitors able to gather data or monitor for deviations by pretending to be a customer?

In the case of personalised pricing, to what extent can customers hide information that firms are using to set personalised prices? Alternatively, could consumer groups, price comparison services, regulators, or the government develop a tool which allows consumers to compare the price that they are quoted with a price based on a 'clean' profile?

- (d) **Replicating studies using UK data** – a final approach could be to replicate the methodology in Chen et al. (2016)⁷³ to assess the prevalence of vendors using pricing algorithms. Similarly, it would also be possible to carry out further research based on the methodologies in Hannak et al. (2014) and Mikians et al. (2012) to conduct a more in-depth and conclusive examination of the extent of online personalised pricing and search discrimination in the UK.

⁷³ Chen et al. (2016), '[An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace](#)', *Proceedings of the 25th International Conference on World Wide Web*, pp1339-1349.

Annex 1: Testing for evidence of personalised pricing

Summary

1. We tested several leading retailers' websites in October 2017 to determine if there was any evidence of personalised pricing. No direct examples of personalised pricing were observed apart from advertised discounts for members. However, we did observe examples of different consumers being shown different search results.

Background

What is personalised pricing?

1. Personalised pricing is the charging of different prices to different customers for the same product. These are not price differences caused by quantity discounts or related to the costs of serving that customer (such as local customers incurring lower delivery fees). The personalisation should be related to the willingness to pay and the price elasticity of the customer.
2. Personalisation could occur in many circumstances such as in direct negotiation in shops or when a customer suggests they are considering switching to another provider. Here we are looking at whether businesses can gather and use data to help them determine the willingness to pay without engaging in direct negotiation and altering the headline or advertised price.

How might algorithms or Big Data make personalised pricing more likely?

3. The use of algorithms in online markets makes this sort of personalised pricing far more effective in terms of how precisely willingness to pay can be identified and the costs of implementing personalised pricing. Online retailers can therefore use personalised pricing algorithms to seek to increase profits.
4. Further, online markets also allow customers to be dealt with individually, so that they can be offered deals which are unknown to other customers.
5. Algorithms may have the capability to personalise prices and to determine what customers' features are associated with a higher willingness to pay. Personalised pricing is more likely with Big Data. Data must be collected on all potential customers and the more data that is gathered the more likely a meaningful relationship can be found that can be used for personalising pricing.

Rationale for new tests

6. The OFT considered personalised pricing in 2013 and did not find evidence of this in the UK.⁷⁴ The OFT noted some examples of newspaper and policy articles (often based on activity in the USA) on personalised pricing, but despite using these to target its approach to personalised pricing, found there was no evidence of this practice in the UK. The OFT doubted that UK consumers were about to face higher prices because of their identity.
7. AK Wien (the Austrian Ministry of Labour, which also has some consumer protection functions) tested for personalised pricing in 2016 and did not find any evidence. However, personalised pricing was found to at least some extent in most of the retailers they tested in March 2017. However, it is not clear exactly what approach was taken to testing for personalised pricing in this study. In particular, it appears that at least some of the price differences they found were dynamic pricing rather than personalised pricing.

Experiment overview

Purpose

8. We wanted to test the extent of personalised pricing in online retail. The aim of this work was to obtain a prima facie indication of whether personalised pricing, as opposed to dynamic pricing, exists in the UK. If so, further research would be needed to assess in greater detail why price variations occurred. We firstly focussed on websites and tests that were most likely to give rise to price differences. We observed prices with multiple users at the exact same time.⁷⁵
9. The study was not designed to look at cross-country differences as there may be good reasons for price differences, such as costs of dealing with a different legal regime in that country or different delivery costs. We considered the possibility of automating some of the tests, but given the range of retailers and products this was not possible.

⁷⁴ Office of Fair Trading (2012), [Personalised Pricing](#).

⁷⁵ Many of the websites change prices regularly (in fast moving markets such as travel prices can change very quickly especially if a particular hotel has received a large order or a flight has been reserved, or it gets close to departure date). Thus, if an item is viewed at different times it may be prudent to refresh the page that was loaded earliest to see if any price discrepancies are unique to the test situation or if the site has globally changed its prices.

Parameters

10. We checked price variances in response to:
 - a. Operating System: past research suggested that a difference in operating systems (e.g. iOS and Android) may cause a difference in final prices. This test could be performed by looking at the results obtained using a Windows operating system and a Mac operating system.
 - b. Logged in vs normal search: Logging in to the website, and therefore revealing the customer's identity, may result in a different price.
 - c. Direct vs Indirect search: Accessing a website directly or via a digital comparison tool or affiliate may affect the price. We did not look at personally tailored products like insurance, but at goods where we expect a single price to be offered. We thus looked at both comparison tools where consumers go solely to choose between the end retailer sites, and cashback sites (such as Quidco and Top Cashback, a type of reward website that pays its members a percentage of money earned when they purchase goods and services via its affiliate links) where consumers assume they are getting a better deal on any site than the other customers that use that site because they are getting the cashback on top. To ensure cashback sites are doing what they say they are, the consumer should pay the same 'headline' rate, whether or not they are using a cashback site. To reduce complexity, we focussed on two cashback websites for each product: Topcashback was the first one used but checks were done using Quidco, froggybank, or comparison sites such as Kayak or Pricerunner.

11. We considered checking price variations in response to:
 - a. Geographic location within the UK: different geographic location may affect the prices offered to users. Ideally, to do this one would need to be in various locations as the pricing effect would normally be implemented by the computer detecting where the user is based using its IP address. Alternatively, this test might be implemented by logging in with a profile that includes a rich or a poor postcode. Unfortunately, the CMA does not have proxy servers or IP addresses that geolocate to anywhere outside London, so we were unable to test easily test this feature. In any event, geographic location is increasingly difficult for retailers to detect and exploit, especially as much browsing and buying now takes place on smartphones, which obtain new IP addresses as users move across areas with different cell towers, but also because IP geolocation may be quite inaccurate even on desktop computers, depending on how Internet

Service Providers assign its dynamic IP addresses.⁷⁶ Additionally, price variation for physical goods (even for prices before delivery) could be due to different costs (e.g. for local warehouse). We therefore decided not to check for price differences due to different locations within the UK, as overall this test was too complex for us to be properly performed.

- b. Purchase history: if past purchasing habits or patterns are observable, firms may use this information to personalise final prices. Ideally we would test whether loyal customers (i.e. customers likely to visit and buy from the same site) were charged more for their loyalty and apparent lack of search than a less loyal one (e.g. a customer profile for someone that often searched and compared options on competing sites). It could also be that after the retailer's pricing algorithm learns that a certain customer profile is connected to someone with a high willingness to pay they start to charge that customer more. Algorithms could easily consider the number of times that customer previously viewed the product or information about products which were bought or viewed by people with similar profiles to perfectly discriminate and set a price equal to the maximum amount a consumer is willing to pay. Overall, however, this would be too expensive for the CMA to test (requiring customer profiles to be built up, and potentially requiring us to either buy products, or to book rooms/tickets and cancel them, which may have research ethics implications). An alternative approach would be to get the public to send us screenshots of different prices for the same trip at (almost) the same time. However, this was too difficult to do within the CMA internet lab given the short timeframe available.

Retailers

12. In terms of retailers to be included we looked at the findings of the AK Wien study and included in some of the biggest retailers in the UK. The largest single price difference that AK Wien found was for Opodo (prices of flights had a variance apparently due to personalisation of at least 6% and up to 40%), so this firm was included. Booking.com was also included for similar reasons.⁷⁷ Other travel sites, that we included were Expedia, and Ryanair (to have an individual airline). We also included Amazon given how important they are in online commerce. Furthermore, we considered major retailers Asda and Tesco who are technically advanced and have been looked at in

⁷⁶ IP address geolocation accuracy is quite high at a country level, but becomes much more hit-or-miss at a regional or city level. See <https://www.iplocation.net/geolocation-accuracy>.

⁷⁷ Showed a 0-11% price difference on hotel rooms.

relation to a Which? complaint on in-store prices but not for their online pricing. The US examples of personalised pricing feature shoplet.com and Staples: we included Staples to take account of these experiences.⁷⁸ Leading European multichannel retailers include Apple, H&M, Zara, Boots, Ikea, and Nike. We chose Apple, and Zara from this list as they are amongst the more technically advanced of the group and have a broad customer base including people with a high willingness to pay who could be adversely affected by personalised pricing.

13. We tested 30 products across ten vendor websites: i) Opodo; ii) Booking.com; iii) Ryanair; iv) Expedia; v) Amazon; vi) Staples; vii) Asda; viii) Tesco; ix) Apple; and x) Zara.

Products

14. AK Wien tested 33 products across 7 websites (3 to 9 products per site). We started with a small exercise of 3 products per site across these 10 retailers (with products sometimes different for each attribute test). Similar to the AK Wien study, we outlined a set of possible products for each retailer before starting but did not specify this rigidly (which may have caused the test to fail, for instance if we chose a product that was out of stock).⁷⁹ In our view, it did not matter which products we compared, so long as the same product compared in each test. We avoided products that were too well known or heavily advertised (since retailers may avoid personalising these as it is clear to customers the price they should be paying).
15. It was important to clear the cookies collected on the web browsing session regularly. This reduced the chance that the retailers would realise they were being tested by having the same two computers repeatedly viewing the same products at the same time.
16. We conducted at least 90 initial tests (10 firms and 3 products for each of three variables), and then re-tested any examples from the 90 where we found price differences to confirm our results and see if these differences appeared random or not.

⁷⁸ The UK website adjusts to show prices with and without VAT.

⁷⁹ We did not search for particular products on CMA machines before accessing the lab to avoid sites having a log of connections from CMA machines to a particular product page.

Results

17. We found very little evidence of personalised pricing but there did seem to be differences in search results that at times were substantial and may have led to different consumer choices.

Operating system (Windows v Mac) comparison:

18. When comparing OS's for Amazon we found a different order of search results. This could reflect price steering (i.e. encouraging some customers to spend more by making higher priced options more prominent), but appears mostly down to the default or selected product category. Opening up Amazon from the two different computers sometimes resulted in a different default search category being applied to the same search terms. We ensured that the search terms were the same on both computers but did not explicitly alter the product category. However, even when the default category was the same for both searches, on some searches there were still some differences in the number or the order of results.

DCT/Cashback link v direct:

19. Amazon did not give any cashback so cashback websites could not be tested for Amazon. There was no cashback offered for Zara so we checked ASOS as an alternative. Booking.com did not show up on comparison sites like Kayak, so was not tested for this. A 128-tool attachment set was advertised at a lower price on the Pricerunner website than the actual price on Amazon but the end price you would pay would still be the same as if you bought it directly without going via Pricerunner (this was probably an error of Pricerunner not personalisation by the retailer). For Asda, Pricerunner displayed a higher Asda price than Asda did but only because Pricerunner included the delivery fee whereas Asda does not but says that customers can collect in store for free.
20. When comparing hotels on Expedia the ranking (order of the hotels in the search results) was different for direct access and when coming via cashback site Quidco. With car hire using froggybank the advertised price is £1 higher on the cashback 'froggybank' link than the direct search, but only because froggybank rounds every number (even 10p over) up to the pound. The final pay page was the same whether you arrived there via froggybank or directly but the link between the pages was different with the direct search having an extra buying step (e.g. to choose to add a car seat). There were other differences in page display/graphics e.g. pages generated by cashback sites often listed the amount of cashback next to the price but also claimed far

more often that a price was a discount compared to normal rates (directly clicking on the search just displayed the current price without suggesting there had been a 'was' price).

21. On one of the websites (Opodo) with one of the cashback sites (Topcashback) we did get a result where the same car hire options for the same booking were £11 (9%) higher via the cashback link than when going direct. However, this result could not be replicated either for the same booking or a different one on that day or on the other days that we did the testing (after cookies had been removed). Thus, we can only assume this was an anomaly. This was only found on the car-hire (5.8% cashback) and not on flights or hotels where cashback was below 2%. There were no price differences from the affiliate links for the other sites.

Comparing logged in prices to direct access:

22. Expedia explicitly advertises that members get special discounts of 10% on particular hotels and we confirmed that although there are no member discounts on car hire or flights they do give up to 11% off on some hotels (which are marked with a yellow reduction). However, this practice is also advertised as available and mentioned on the non-logged in site. For car hire the logged in search appeared to default to searching the whole city rather than just the airport. For the hotel search the second result was a more expensive choice for the logged in customers.
23. Several firms showed different search results to the logged in and browsing customer, including:
 - a. On Asda, for the Lego search the third result was a more expensive choice when not logged in.
 - b. On Zara, there was a difference in one of the search results (but not the prices).
 - c. On Booking.com the edreams price was higher (£67 when not logged in rather than £61) but the price was not the lowest price found so should not affect the customer. The number of results changed substantially with 149 results when not logged in and 307 otherwise.
 - d. On Opodo when searching for hotels the first three results for browsing cost £720, £220, and £300. While for the logged in customer the options displayed cost £570, £1500, and £680. Thus, the logged in customer may end up being persuaded to pay more. When this was tested for a different hotel location the list of search results for the two customers (logged in

and not) were different but there was no real pattern of one customer being shown higher priced options than the other.

- e. Tesco gave slightly fewer results for some searches but not always to the same person. Twice the unidentified browsing customer got more results (usually in the same order) while once it was the logged in customer.
- f. The other retailers gave the same prices and the same order of search results.

Areas for improvement

- 24. Given the wide range of potentially relevant factors (which can interact), it is unlikely that experiments focusing on varying one factor at a time would uncover personalised pricing.
- 25. A more sophisticated method which is likelier to uncover personalised pricing would be to recruit real customers to collect data on prices.⁸⁰ These real customers are likely to have built up a 'profile' or pattern of signals which would allow firms to respond with a personalised price. It is technically possible to conduct a rigorous test which removes any effect of dynamic pricing, by setting up the customers' system so that when they request a price on a webpage, a simultaneous request is made from a 'blank' profile and recorded. Whilst this method would likely uncover evidence on the existence of personalised pricing, it would not reveal what exactly about the customer's profile is used as the basis for personalised pricing.

⁸⁰ This would be a kind of crowd sourcing of evidence – either voluntarily (although the quality of the evidence may be poor) or paid by using the techniques outlined in articles such as: Measuring Price Discrimination and Steering on E-commerce Web Sites, 2014, Hannak et al. http://www.ccs.neu.edu/home/cbw/static/pdf/imc151-hannak.pdf#_ga=2.142287218.1509736259.1512384335-691489406.1512384335. One example that suggests personalised pricing may be happening in the UK is that the shirt retailer TMLewin appear to price their shirts at one price if you navigate directly to their webpage (often £39.95) but if you search for a competitor "Charles Tyrwhitt" and then click on the pay per click TM Lewin link a significantly lower price is displayed (£22.50 or £19.50). Such an example is at this stage just anecdotal.

Table 1: Tabulated Results of Personalised Pricing Tests

| Retailer/test | Operating System | DCT/cashback link | Logged in to site |
|----------------------|---------------------------------|---|--|
| Amazon | Different search results order. | - (no cashback), Amazon was cheaper than Pricerunner but probably due to error in reporting price by Pricerunner. | - |
| Apple | - | | - |
| Asda | - | Pricerunner higher price than Asda but due to including delivery cost (rather than in store collection). | - |
| Booking.com | - | - (not on comparison sites) | Difference in one quote but not the lowest quote, different number of search results |
| Expedia | - | Different search results order. With car hire using froggybank higher price stated due to rounding and different options (car seat). Different graphics and means of displaying discount. | Advertises that members get 10% off at selected hotels. Different search results. |
| Opodo | - | One-off result of more expensive (9%) car hire rates for the Topcashback link but could not be replicated. | Logged in customer appears to get more expensive hotels recommended. |
| Ryanair | - | - | - |
| Staples | - | - | - |
| Tesco | - | - | Slight difference in number of search results |
| Zara | - | - (ASOS was tested) | Slight difference in search results order. |