

ARTIFICIAL INTELLIGENCE AND PRICING

Diego Aparicio and Kanishka Misra

ABSTRACT

As businesses become more sophisticated and welcome new technologies, artificial intelligence (AI)-based methods are increasingly being used for firms' pricing decisions. In this review article, we provide a survey of research in the area of AI and pricing. On the upside, research has shown that algorithms allow companies to achieve unprecedented advantages, including real-time response to demand and supply shocks, personalized pricing, and demand learning. However, recent research has uncovered unforeseen downsides to algorithmic pricing that are important for managers and policy-makers to consider.

Keywords: Artificial intelligence; pricing; personalized pricing; dynamic pricing; pricing algorithm; supracompetitive; price discrimination

1. INTRODUCTION

Advances in technology enable businesses to overcome barriers that pricing managers have faced for decades. For example, “menu costs” associated with updating retail prices have been a common obstacle. Menu costs include the costs of the printing and physical placement of a new price tag every time the price for a given product changes. As a result, setting different prices for different varieties of a product (e.g., different shades of a nail polish) is cost prohibitive for stores. [Anderson, Jaimovich, and Simester \(2015\)](#) suggest that because of labor costs, price changes on products average only 100 SKUs per day, for five days a week (out of about 20,000 SKUs). Another barrier stems from firms' prevailing managerial practices ([Aparicio & Simester, 2022](#)). In fact, [Dellavigna and Gentzkow \(2019\)](#) report nearly uniform prices (across stores) for most US food, drug store, and mass merchandise chains – even though individual retail stores face different demand conditions.

Artificial Intelligence in Marketing

Review of Marketing Research, Volume 20, 103–124

Copyright © 2023 Diego Aparicio and Kanishka Misra

Published under exclusive licence by Emerald Publishing Limited

ISSN: 1548-6435/doi:10.1108/S1548-643520230000020005

Dellavigna and Gentzkow (2019) describe this phenomenon as managerial inertia – a combination of agency frictions and behavioral management factors. Hitsch, Hortaçsu, and Lin (2021) complement this evidence with a statistical barrier: given the available historical data, credibly distinguishing price elasticities across stores is hard for managers. This challenge, in turn, limits the feasibility of price discrimination across stores.

As stores digitize, some of these barriers disappear. Today, an e-commerce pricing manager can change the price of a product with the press of a button. The menu cost is near zero, and indeed we observe firms updating prices in very short time intervals, e.g., every 15 minutes (Chen, Mislove, & Wilson, 2016; White, 2012). In physical retail stores, the increased use of digital shelf labels (Hansen, Misra, & Sanders, 2021a; Stamatopoulos, Bassamboo, & Moreno, 2021) also facilitates such changes. However, these digital innovations present new “big data” challenges for pricing managers. Without the constraints of physical displays, e-commerce retailers provide consumers with much larger product assortments.¹ In addition, in e-commerce, firms can observe rich data about consumer behavior (e.g., search, clicks, and purchase decisions) and competitive pricing in real time. These unprecedented changes in scale (Baker, Kiewell, & Winkler, 2014) require heavy investments in computing resources. Even with these rich data and sufficient computing power, how can managers/analysts set real-time prices across millions of individual products?

Artificial intelligence (AI) provides a solution: A machine receives a set of instructions to *automate* the pricing decision. The purpose of this paper is to survey current research in the area of AI and pricing. In Section 2, we consider research that provides evidence and insights for firms that are implementing AI for pricing. We discuss the potential advantages: automatic pass-through from cost shocks to prices (e.g., swings in commodities); price discrimination (i.e., a firm can charge different users a different price); and price discovery (i.e., experiments to discover the profit-maximizing price). In Section 3, we consider a longer term and broader view of how the increased use of AI models affects firms, consumers, and public policy-makers. Here, we review studies suggesting that firms can change consumer preferences with many price changes, as well as studies examining the welfare implications of price discrimination. For example, as the prevalence of AI adoption increases, it raises the competitive bar and causes rival firms to race and react in real time to the other firms’ prices. Section 4 concludes with some key learnings and avenues for further research.

2. FIRMS IMPLEMENTING AI FOR PRICING

AI models provide the backbone infrastructure that allows firms to *automate* their pricing decisions. By “automation,” we mean that a computer algorithm determines the current price. This algorithm could be as straightforward as defining some set of conditions, such that if those conditions are met, the price is automatically updated. Or the algorithm could be a complex learning model that

analyzes real-time data to set current prices. In most implementations, the pricing algorithm operates as a computation mechanism using inputs such as pricing rules (Proserpio et al., 2020).

For historical background, airlines were among the first firms to use AI to automate their pricing decisions (Borenstein & Rose, 1994; McAfee & Te Velde, 2006). Setting the price for a flight is inherently a complex, dynamic problem because it depends on many elements that tend to exhibit abrupt shifts from day to day. Airlines are subject to sudden changes both in demand (customers' booking patterns) and in capacity (seat availability). As a result, research reports a large degree of price variation across days in airline tickets, depending on the time of purchase (relative to departure), competition on a particular route, and the number of seats sold.

AI technology dramatically affects the price patterns observed in many markets. In some cases, the price variability can be quite striking. Fig. 1 shows two examples. Panel A collects the prices for an identical ride provided by UberX, going from Boston's Museum of Fine Arts to Boston's Celtics Stadium (Aparicio & Rigobon, 2021). Over the course of 7 days, there were 156 distinct prices, ranging from \$8.23 to \$14.38; and in some cases, prices varied by as little as 1 cent (e.g., \$9.52, \$9.53, \$9.54, \$9.55). The price variation (defined as the ratio of standard deviation to average price) is estimated at 9%. This price dispersion, while high, is not unprecedented; for example, using weekly scanner data from offline stores, Hitsch et al. (2021) show that Tide laundry detergent's price varied between \$5 and \$8 between 2008 and 2012. What is different with AI then? The answer is the time frame in which these price changes are observed. The variation in Uber's prices is observed in just a few days, while the variation in the supermarket's Tide prices is observed over a few years! Panel B in Fig. 1 offers another illustration of the degree of AI-induced price variation. It shows the prices in a US online grocery platform for several products, including cereal, painkillers, personal care, and snacks (Aparicio, Eckles, & Kumar, 2022). This online platform implemented an AI algorithm to set prices. The charts reveal the remarkable variability of prices in periods when the retailer utilizes price algorithms.

The range of businesses that use AI for pricing is wide and growing. What are the advantages of AI pricing? In this chapter, we structure the advantages into three parts:

- Dynamic Pricing: Real-time Swings in Demand and Supply
- Personalized Pricing: Price Discrimination
- Price Experimentation: Demand Learning

2.1 Dynamic Pricing: Real-Time Swings in Demand and Supply

Historically, pricing automation was driven by the need to account for real-time changes in demand and supply. Consider the example of airline tickets. Airlines use AI to automate the pricing decision because seat occupancy and proximity to

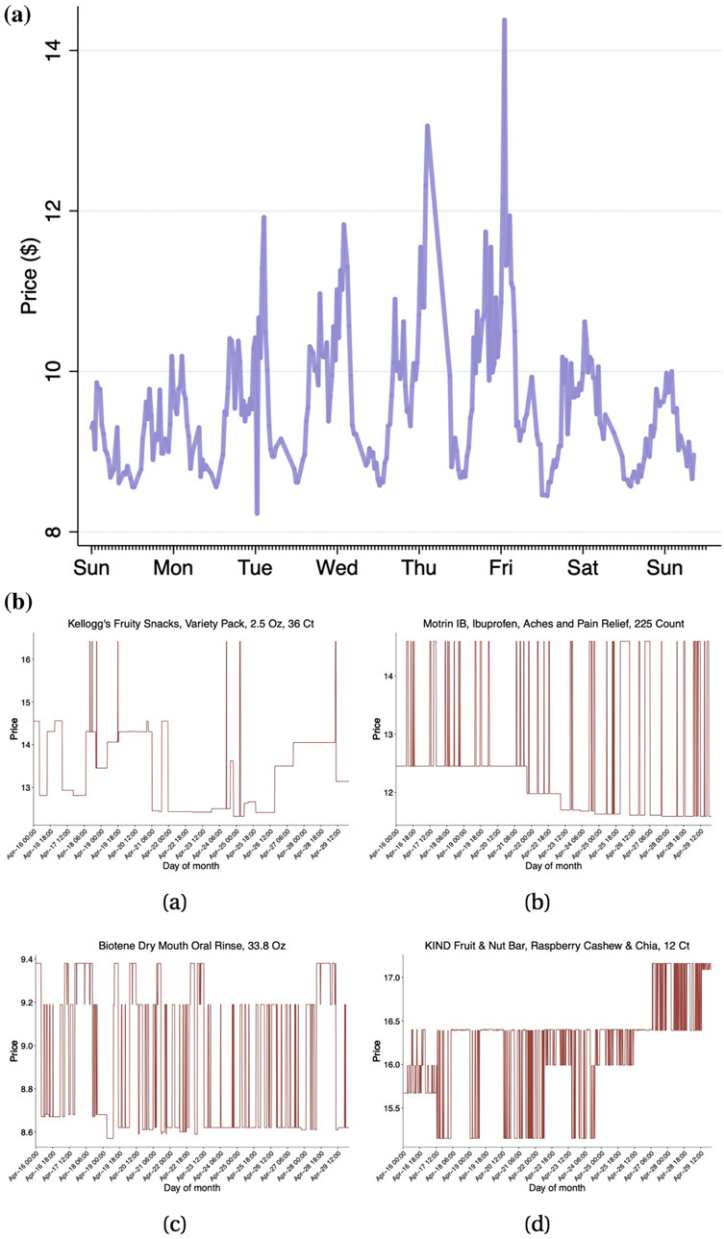


Fig. 1. Examples of Algorithmic Pricing: Uber and Online Groceries. (a) Prices charged for a UberX ride between two points: Boston’s Museum of Fine Arts and Celtics Stadium. Data were collected for 1 week in 30-minute timestamps. (b) Algorithmic pricing in four products sold by US online platform. *Source:* Extracted from Aparicio and Rigobon (2021), Aparicio et al. (2022).

departure vary continuously. Therefore, the ticket price a consumer observes in the morning can be different from the price she will observe later in the afternoon. Here, we consider the use of AI models to automate this process of inter-temporal pricing.

Ride-hailing platforms offer a stylized, well-recognized case of AI models for intertemporal pricing. Platforms like Uber and Lyft use “surge pricing” algorithms to respond in real time to spikes and drops in drivers and consumers (Allon, Cohen, & Sinchaisri, 2018; Castillo, 2020; Hall, Kendrick, & Nosko, 2015). The objective of surge pricing is simple: Increase prices when demand is higher than supply and reduce prices when supply is higher than demand – the principle taught in Microeconomics 101. Consider Uber’s prices in Panel A of Fig. 1. Can fluctuations in demand (riders) or supply (drivers) explain such disparate prices? Castillo (2020) studies data from the city of Houston and suggests that the answer is yes. The study documents remarkable, high-frequency swings in demand and supply at Uber across the hours within a day, and across the days of the week. This imbalance creates enough price differences to allow for structural estimates of agent primitives. The analyses indicate that, on the demand side, consumers assign a high value to time (about \$2 per minute); and on the supply side, drivers have incentives to move to areas with high surge pricing.

The short-term rental market is another environment characterized by abrupt and rapid swings in supply or demand. Huang (2021) finds evidence suggesting that Airbnb sellers (hosts) face cognitive constraints that prevent them from setting the optimal price (a night of stay). The intuition for why sellers might have limited bandwidth is that, similar to airlines or hotels, the pricing problem is difficult for two reasons: (1) some nights face extraordinary demand (e.g., a concert, a long weekend) and (2) the number of nights is capacity-constrained and perishable (i.e., not renting is an opportunity cost). Perhaps tapping into this opportunity to streamline hosts’ rental prices, Airbnb offers an AI solution – “smart pricing” – that sellers can use to automate the pricing decision. However, while this interface eases sellers’ barriers to changing their prices, the prices decided by the “smart pricing” tool may not be profit-maximizing for the seller. Instead, they may be profit-maximizing for the platform. Consequently, some sellers resist adopting this technology.

Airbnb’s “smart pricing” tool is interesting because it illustrates a potential incentives-related problem between the sellers’ pricing problem and the platform’s pricing problem. Also interesting is that it showcases how mainstream users are finding AI models for pricing easier to adopt. A few years ago, AI pricing was primarily exclusive to big, techsavvy players; now, numerous software companies provide AI-based repricing solutions at a relatively low cost (Musolf, 2021). As these types of solutions become even more accessible to small users, future research needs to examine its effects and whether they lead to a more level playing field.

Competition. Businesses increasingly are motivated to use AI for pricing to be more competitive. Competition intensifies especially when strategic behaviors emerge between sellers of the same product. For instance, imagine that

consumers are very price-sensitive and seek the cheapest alternative before completing a purchase; for sellers of the same product (e.g., a used book, a concert ticket), matching a competitor's price can be critical. Alternatively, a firm may be unsure about demand for its product (and the price it should set) and may believe that other firms' prices contain useful information. Or a firm may lack the experience or time to update the price to changing market conditions and may instead prefer to "follow" another firm. Regardless of the microfoundation, AI is an ideal tool because it allows users to construct a set of instructions that automatically change the price on the basis of the competitors' price. The set of instructions can be as simple as: (1) track competitor prices and (2) match the lowest competitor price. Experienced platforms are likely to use a more advanced set of instructions; for example, they may input competitor prices to re-optimize their pricing algorithms, based on expected demand.

Merchants that sell the exact same product on the same platform may be fiercely competitive. Intuitively, consumers that navigate a platform of sellers do not need to "search" for the lowest price: The platform provides this information. For example, Amazon's marketplace (and other platforms) reveals the cheapest alternative merchant for Dan Brown's *Da Vinci Code*, or a pair of Adidas shoes, or for a TV monitor. The early work of [Chen et al. \(2016\)](#) characterizes the existence of AI pricing between competing sellers in Amazon's marketplace (including Amazon itself as a seller). The authors report a number of insightful facts, such as the frequencies and dynamics of price changes, and how/when an individual seller's use of algorithmic pricing increases the seller's success. Figures 14–17 in [Chen et al. \(2016\)](#) show their evidence of sellers' automatic price setting – in some cases, price-matching the lowest price; in other cases, as that of Amazon, setting a constant premium price over the available sellers. An interesting finding of this study is that charging the lowest price is highly predictive of which product is selected in the "Buy Box." Why is this important? To the extent that a large proportion of platform sales occur through the "Buy Box," then adopting AI can be crucial in sellers' attempts to outbid rivals and become the default seller of choice.

Price comparisons are not limited to competing sellers in the same marketplace (i.e., website). Retailers often use competitive price trackers to compare their prices with those of their competitors. As a visual example of price-matching, consider Panel B of [Fig. 2](#), obtained from [Aparicio, Metzman, and Rigobon \(2021\)](#). The graph shows evidence that Amazon and Walmart, the two online grocery e-commerce giants in the United States, match each other's price. Although online demand for ice cream arguably may not be very sensitive to a few cents' difference, other aspects of the technology's use are worth attending to. In particular: (1) AI is used to target a specific variety of ice cream (same exact product) and (2) Amazon and Walmart carefully monitor each other's price and occasionally price-match within hours and for the same delivery ZIP code. This example also showcases the vast amount of data that platforms amass by the hour: prices for each SKU, for each delivery ZIP code, across hours, across competing sellers.




Fig. 2. Algorithmic Price-Matching. Amazon's and Walmart's Price of Ben & Jerry's Chocolate Fudge Brownie for the Same Delivery Zipcode. *Source:* Extracted from Aparicio et al. (2021).

Another study that sheds light on the connection between AI pricing and competition is [Brown and Mackay \(2020\)](#). Using a rich and high-frequency dataset of over-the-counter allergy drugs from online retailers, the authors report substantial heterogeneity in the retailers' pricing technology for the same product. Some retailers change prices within a day, and other retailers change prices only once per week. How do these disparities in their pricing practices affect the overall market for consumers, and how do they affect the retailers' incentives to compete with each other? Again, the authors find that retailers are mindful of competitor prices and leverage their pricing technology to undercut rivals' prices. They show that retailers that have *faster* AI technology (defined as 1–2 hours between price changes vs. 160 hours between price changes) are able to undercut a rival's price more rapidly, resulting in prices that are set about 30% lower. This work further suggests that AI pricing algorithms may lead to non-Bertrand Nash price competition; we will discuss this in Section 3.3.

AI applications for pricing reviewed in this section have a common feature: AI is used to change prices very frequently. Ride-hailing prices, grocery prices, price-matching, and undercutting a rival are all characterized by *high-frequency price changes*. Consistent with these trends, [Cavallo \(2018\)](#) reports a steady increase in the number of price changes at online retailers as they adopt

algorithmic pricing. Indeed, this fact tends to be exploited to identify the use of algorithmic pricing. For example, [Chen et al. \(2016\)](#) define algorithmic sellers based on two factors: (1) the number of price changes and (2) the correlation of prices to the lowest of Amazon's price; [Assad, Clark, Ershov, and Xu \(2020\)](#)'s study of the German gasoline market identifies algorithmic prices based on a trend break in the number and timing of price changes; and similarly, [Aparicio et al. \(2022\)](#) exploit periods with and without intense price changes to infer when the platform is using algorithmic pricing. However, we stress that, while frequent price changes are a notable feature of AI pricing, they are *not* the only feature.

2.2 Personalized Pricing: Price Discrimination

Price discrimination can be a useful tool in the economics of pricing, allowing firms to charge different prices to different sets of consumers.² However, in many cases, firms are not willing or able to charge a personalized price. One reasonable concern is that the media or the public in general might be antagonistic if they know that their specific "user identifier" is being targeted and charged a different price for the same exact product. Amazon provides an exemplary illustration: As recounted in [The Atlantic \(2017\)](#): "In 2000, some people thought Amazon was doing this when customers noticed they were being charged different prices for the same DVDs. Amazon denied it. This was the result of a random price test, CEO Jeff Bezos explained in a news release. *We've never tested and we never will test prices based on customer demographics.*"

The marketing literature has discussed the advantages of using historical consumer data to achieve first-degree price-discrimination (i.e., personalized prices at the individual level). Studies in this area simulate the benefits of optimized targeted coupons and report notable profit gains (e.g., [Gabel & Timoshenko, 2021](#); [Pancras & Sudhir, 2007](#); [Rossi, Mcculloch, & Allenby, 1996](#); [Zhang & Wedel, 2009](#)). Studying the implications of price targeting, [Smith, Seiler, and Aggarwal \(2021\)](#) show (1) a large variation in predicted benefits based on their assumed model of demand³ and (2) the importance of accurate purchase history data (vs. demographics data).

When customization at the individual level is not feasible, firms may "cluster" prices at a broader aggregation level: Stores or consumers in different geographic areas are charged different prices. Studies that have examined geographical price discrimination in brick and mortar retail settings ([Adams & Williams, 2019](#); [Chintagunta, Dubé, & Singh, 2003](#); [Duan & Mela, 2009](#); [Ellickson & Misra, 2008](#); [Hoch, Kim, Montgomery, & Rossi, 1995](#); [Li, Gordon, & Netzer, 2018](#); [Zhang & Krishnamurthi, 2004](#)) (1) provide evidence for different pricing strategies across stores and (2) simulate the benefits of customized pricing for each store. For instance, [Chintagunta et al. \(2003\)](#) study the implications of price discrimination at different geographic levels (chain, zones of stores, stores), and they estimate that store level pricing can increase profits for the retailer by about 10%–15%.

Despite the potential benefits of heterogeneous pricing strategies, [Dellavigna and Gentzkow \(2019\)](#) document that prices within a retail chain tend to be fairly

uniform (robust to food, drugstore, and mass-merchandise chains). They estimate that this lack of cross-sectional (geographical) price discrimination costs the median retailer about \$16 million in annual profits (about 1.6% of revenue). The reasons for this uniformity and the related cost include both managerial practices and statistical limitations of insufficient price variation (Hitsch et al., 2021). This lack of price discrimination is not unique to a brick and mortar retail context. In many other industries, we observe that firms do not price-discriminate across their product portfolio. Notable examples include prices at movie theaters (Orbach & Einav, 2007), where all movies – regardless of popularity – have the same price; rental cars (Cho & Rust, 2010), where all cars – regardless of mileage – have the same price; online music (Shiller & Waldfogel, 2011), where prior to 2009, Apple’s iTunes charged 99 cents per song across the board; and fashion (Aparicio & Rigobon, 2021), where products that have different colors, models, or fabrics have an identical price.

AI price algorithms overcome barriers of uniform pricing. Aparicio et al. (2021) collect novel data for hourly prices across delivery ZIP codes for the leading online grocery retailers in the United States. They provide evidence that most online grocers, but especially Amazon and Walmart, personalize prices at the delivery ZIP code. Thus, two customers looking to buy Oreos at 9:00 p.m., one located in Miami and another located in Chicago, will be charged different prices. Fig. 3 shows visual examples of price personalization based on the delivery ZIP code. Note that even if customers somehow realize these price differences, they will not be able to “arbitrage” them – that is, a customer is unlikely to have their groceries delivered to another area to save the difference. The study also shows that algorithms that price-discriminate (across locations) are also more likely to price-change (across time). Therefore, as firms adopt more AI tools for pricing, we might start to see more price discrimination as well.

We do note that the practice of geographical price discrimination is not homogeneous across all online retailers. Using data from 2008 to 2013, Cavallo, Neiman, and Rigobon (2014) report that Walgreens and Walmart were among the few largest retailers to use location-based price discrimination. Cavallo (2017) collected data from 56 retail chains in 10 countries and documents that the online price matches the offline (physical store) price about 72% of the time. Moreover, Cavallo (2018) considers the role of competition in this lack of geographical price discrimination, particularly for durable goods, and suggests that if a competitor (Amazon.com in these data) does not price-discriminate, then another firm (Walmart in these data) is less likely to do so as well.

In the extreme case, online prices may be personalized for every consumer. Perhaps one of the most compelling studies on price personalization is Dubé and Misra (2021). The authors collaborated with Ziprecruiter.com⁴ and implemented a field experiment to personalize prices for new customers. Ziprecruiter.com’s status-quo was a price of \$99 per month for all consumers; the experiment randomized prices from \$19 to \$399 per month. The authors estimated a heterogeneous demand model and optimize the personalized prices based on an observable set of 133 characteristics variables of a potential customer.⁵ The study finds economically large effects from personalized price: 19% increase in expected

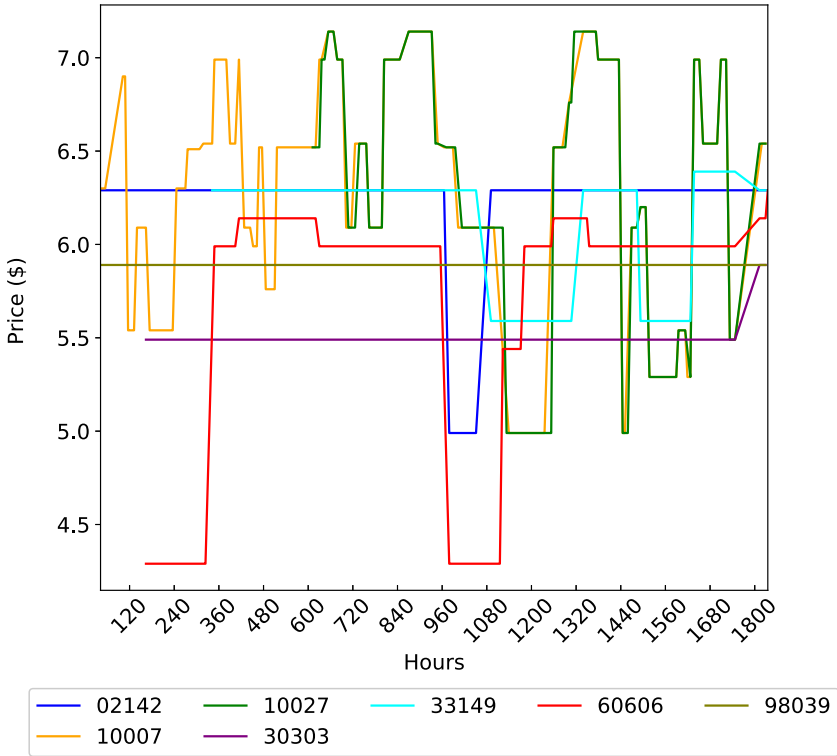


Fig. 3. AmazonFresh's Prices for a 12-Pack Diet Coke, Depending on the Delivery Zipcode. *Source:* Extracted from Aparicio et al. (2021).

profits relative to the optimal uniform price. This chapter raises a potential concern when implementing model-based price discrimination: the highest personalized price is \$6,292, which is more than 15 times the maximum price (\$399) in the experiment. Even limiting the maximum personalized price, they estimate a 8% increase in revenue over uniform pricing. Importantly, they verify this estimate with a second field experiment where they randomly assign potential consumers to uniform and personalized pricing conditions.

The use of AI for personalized prices is not exclusive to online markets. [Karlinsky-Shichor and Netzer \(2019\)](#) assess historical pricing data to recommend personalized prices that sales people can use in their negotiations. Conducting a randomized field experiment, the authors show that adding the AI-based price recommendation to the customer relationship management system can increase profitability by 10%. In a similar vein, [Cui, Li, and Zhang \(2022\)](#) study the role of AI in procurement in an online B2B platform that connects buyers and suppliers. The authors provide evidence showing that the quoted wholesale prices react to the use of AI – quoted prices exhibit a premium for chatbot buyers, but the effect

is moderated if the buyer signals the use of AI to screen suppliers. This blended approach of using AI models to enhance managerial decision-making is discussed in Proserpio et al. (2020) and can provide a promising avenue for offline firms to start using algorithmic pricing solutions.

2.3 Price Experimentation: Demand Learning

To implement a pricing strategy, a manager has to answer the most fundamental question: What is the optimal price for my product? It is not an easy question in general. But it is even more difficult when the firm lacks data on the demand (i.e., a new product, a change in the competitive landscape, or a change in policy).⁶ A proactive solution is to *experiment* with various prices. There are different kinds of experiments. Aparicio et al. (2021) provide evidence that online grocers explore the price grid, an indication that price algorithms are regularly running small experiments to test many distinct prices. Similarly, firms can make small, randomized price changes to better estimate price elasticities (Fisher, Gallino, & Li, 2018). In fashion, firms can introduce products at different price buckets and later estimate the best bucket (Aparicio & Rigobon, 2021), or firms can dynamically adjust markdown prices based on demand signals (Cachon & Swinney, 2011). A novel illustration of price experimentation is Dubé and Misra (2021). In collaboration with a B2B online platform, the authors were able to *randomize* the prices charged for new customers. The field experiment covered real purchases, as well as a vast range of prices: customers were randomized to prices between \$19 and \$399 (with a baseline price of \$99).

Instead of considering learning and earning as two distinct phases, AI models frame this as a dynamic optimization problem with the goal of maximizing *earning while learning*. The computer science literature considers this problem in the class of reinforcement learning (Sutton & Barto, 1998). Within this category, *multi-armed bandit* (MAB) methods provide algorithms for describing the experimentation process (e.g., Auer, 2002; Gittins, Glazebrook, & Weber, 2011). In the application of MAB methods to pricing (see Kleinberg & Leighton, 2014; Misra, Schwartz, & Abernethy, 2019), these methods consider a discrete set of possible decisions (prices), each with a stable but unknown profit. The pricing algorithm sets prices to balance current earning of profits and learning about demand with future profits.

To see the advantage of this approach, consider Fig. 4 obtained from Misra et al. (2019). In Panel A, we see that a balanced field experiment (left), charges all prices with the same frequency. The key difference with a bandit (right) is that the algorithm incorporates additional real-time feedback to update the prices charged in real-time. The bandit's objective of earning while learning leads to higher profits (reduced experimentation costs) – after 1,000 prices, the bandit earns 95% of optimal profits, while the balanced experiment earns only 66% of optimal profits. Panel B shows the information learned about the true profit curve. The balanced experiment (left) learns profits evenly at all price points, while the bandit (right) learns most accurately about profits at near the optimal price

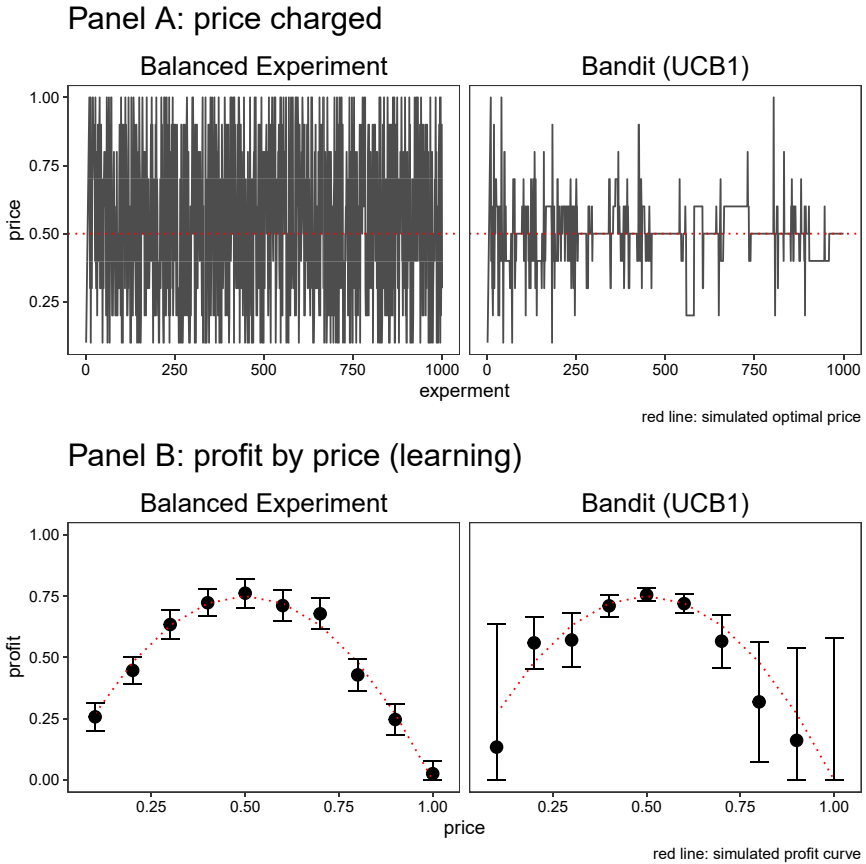


Fig. 4. Comparing a Simulated Balanced Field Experiment to a Multi-armed Bandit. *Panel A. Prices Played:* shows the spot price charged over time. A balanced experiment (left) charges prices with equal probability while the bandit (right) leans in real time and charges the truly optimal price (\$0.50 in this simulation) more often. *Panel B. Profits Learning:* The chart shows the mean and 95% confidence intervals of the learned profit at each price. The red dotted line represents the true profit curve. The balanced experiment learns the profit curve with the same precision at all prices, while the bandit learns the true profit with small confidence intervals around the optimal price (\$0.5 in this simulation) and large confidence intervals at suboptimal prices. *Source:* Extracted from [Misra et al. \(2019\)](#).

(smaller standard errors). This is similar to the notion of “adequate” knowledge in [Aghion, Bolton, Harris, and Jullien \(1991\)](#).

To find the profit maximizing price, a researcher can assume a parametric functional form for demand to impose institutional or economic knowledge. The learning problem then becomes one about learning the parameters of the demand

function. An optimal algorithm to consider in this setting is called Thompson sampling (Thompson, 1933).⁷ Ganti, Sustik, Tran, and Seaman (2018), from Walmart labs, provide an example implementing this algorithm. They consider a simple constant elasticity (single parameter) demand model per product and show that applying Thompson sampling for dynamic pricing can lead to an improvement of per unit sales, particularly when Thompson sampling includes a large fraction of the products sold. Hansen et al. (2021a) consider data from a brick-and-mortar retailer that uses digital shelf labels to set discounts (using Thompson sampling) for expiring perishable products.

An alternative to this setup is to consider using a nonparametric approach, where the AI model learns demand fundamentals. An optimal algorithm is the upper confidence bound (UCB) algorithm (Auer, 2002).⁸ In a pricing application, Misra et al. (2019) add demand learning to the UCB algorithm by enforcing that demand curves must be weakly downward sloping. They show combining economic theory and MAB methods increases the profitability of experiments, compared to balanced field experiments and standard methods from computer science.

A limitation of MAB is that they are single-state models (Sutton & Barto, 1998). A richer framework is one in which the underlying demand state can change endogenously – depending on the price charged (action taken). A nonparametric solution is the Q-learning algorithm from the reinforcement learning literature (Sutton & Barto, 1998; Watkins & Dayan, 1992).⁹ In pricing, a Q-learning algorithm has been implemented in several contexts; for example, see applications in energy (Lu, Hong, & Zhang, 2018), online coupons (Liu, 2021), and perishable retail (Cheng, 2008). These applications differ in the endogenous state considered, in Liu (2021) the endogenous states are defined by the consumers' RFM (Recency Frequency Monetary value) status, while in Cheng (2008) the endogenous state is the inventory of the perishable good. Calvano, Calzolari, Denicolo, and Pastorello (2020) and Klein (2021) consider multiple firms simultaneously using Q-learning to study competitive market outcomes. In these applications, a history of all prices defines the endogenous state and each firm's pricing algorithm can derive the optimal pricing strategy as a function of competitive prices.

3. CONSEQUENCES OF AI FOR PRICING

In this section, we provide an overview of the research highlighting potential consequences of AI, for each of the advantages we discussed earlier.

3.1 Dynamic Pricing

The objective of AI algorithms is to set prices based on real-time primitives (consumer preferences and supply). This idea is not new and the technology to implement such pricing has existed for decades (Seele, Dierksmeier, Hofstetter, & Schultz, 2021). However, dynamic pricing is still not implemented in many markets. A famous example dates back to 1999 when the Coca-Cola Company

tested vending machines that varied price based on the weather.¹⁰ The CEO and chairman, Douglas Ivester, described how the utility for a Coke can increase during sporting events and during summer heat and said “the machine will simply make this process automatic.” The public outcry after these comments was intense, and these machines were never launched!

From a microeconomic perspective, if AI algorithms unintentionally shift primitives, then the derived price may no longer be optimal.¹¹ One mechanism through which this shift could happen is via consumer concerns about fairness. Increasing prices during periods of high demand might inspire consumers to question firms’ fairness (Kahneman, Knetsch, & Thaler, 1986; Rotemberg, 2011). In the Coke example, rival Pepsi’s spokesperson raised these concerns and said, “We believe that machines that raise prices in hot weather exploit consumers who live in warm climates.” In laboratory experiments, Haws and Bearden (2006); Weisstein, Monroe, and Kukar-Kinney (2013) show that participants’ perceptions of companies’ fairness and trustworthiness are affected by dynamic pricing. Further, Feinberg, Krishna, and Zhang (2002) show that these fairness concerns can exist with targeted promotions.

A second mechanism for shifting attitudes toward AI pricing could come through consumers’ learning about their own primitives through their observations of price variation. Aparicio et al. (2022) find that consumers exposed to algorithmic pricing in the context of online grocery shopping become more price-sensitive, compared to similar users buying the same products, in the same time periods, without such price exposure. This finding complements earlier work showing that consumers who receive promotions may become more prone to seek deals in the future (Elberg, Gardete, Macera, & Noton, 2019; Zhang et al., 2020).

The welfare implications of dynamic pricing are nontrivial. Castillo (2020) estimates that Uber’s AI-based surge pricing increases welfare for riders, relative to a benchmark of uniform pricing. The effect is driven by two factors: (1) a better allocation of cars (i.e., during times and in areas of high demand) means that riders benefit from being more likely to find a ride when demand is high; and (2) surge pricing policies allow for lower prices during times of lower demand (vs. uniform pricing). However, in times of crises, a purely automated price strategy also can lead to price gouging (Snyder, 2009). For example, airline prices during Hurricane Irma and prices for personal care items during COVID-19 have received abundant media coverage.¹² During the recent pandemic crisis, the US Congress did introduce a Price Gouging Prevention Act,¹³ which “makes it unlawful for any person to sell or offer for sale a consumer good or service during a public health emergency resulting from COVID-19 (i.e., coronavirus disease 2019) at a price that (1) is unconscionably excessive, and (2) indicates that the seller is using the circumstances related to the public health emergency to increase prices unreasonably.” In 2020, the New York Attorney General did fine three sellers on Amazon for price gouging of hand sanitizers at the start of the crisis.¹⁴ The three sellers were ordered to pay a fine and to reimburse consumers who had purchased these products.

3.2 Personalized Pricing

The White House's Council of Economic Advisors (CEA) ('White House', 2015) addressed the scope of AI pricing algorithms. The CEA cautioned that AI pricing algorithms lower costs of first-degree price discrimination, leading to a transfer of welfare from consumers to firms. Similar questions have been raised in the revenue management literature (Van Der Rest, Sears, Miao, & Wang, 2020). The CEA did report that this concern may be somewhat limited, and no policies were enacted to legally limit the use of personalized pricing. Important to this debate, Dubé and Misra (2021) quantify the welfare gains and losses of personalized prices. Their results indicate that personalized pricing reduces overall customer surplus by 23%, however is welfare enhancing for a majority (over 60%) of consumers (primarily smaller firms) who benefit from a lower price (vs. uniform pricing). They estimate the role of policy restrictions for personalized pricing and find that sometimes the full model generates more consumer surplus than several of the restricted scenarios. Therefore, they conclude: "over-regulation of the types of data firms can use for personalized pricing purposes could exacerbate rather than offset some of the harm to consumers." In other markets, Kallus and Zhou (2021) indicate that personalized pricing in elective vaccines and micro-finance can be welfare enhancing. Future empirical work can consider these welfare implications under different industries and market conditions.

The area of personalized pricing relates to broader questions about possible downsides of algorithmic management, or "algorithmic bias" (e.g., Hajian, Bonchi, & Castillo, 2016; Lambrecht & Tucker, 2019). In this literature, researchers study potential discrimination among users based on, for example, race, age, and sexual orientation. As price algorithms learn and keep fine-tuning their possibilities to become more and more targeted, they may become vulnerable to learning and exploiting such potential biases.

In the Airbnb context, Zhang, Mehta, Singh, and Srinivasan (2021) study the extent to which the "smart pricing" tool indirectly reduced or widened the racial gap among hosts (sellers) on the platform. The question is relevant in terms of both policy and economic outcomes because technologies that automate prices may input features, or learn from demand curves, that incorporate racial economic inequality and, unintentionally, set prices that exacerbate this inequality (Cowgill & Tucker, 2019; Kleinberg, Ludwig, Mullainathan, & Rambachan, 2018). Zhang et al. (2021) report quasi-experimental evidence that the algorithm benefited black adopters more than white adopters, which in turn decreased the racial revenue gap by more than 70%. The algorithm set similar prices for equivalent properties owned by black and white hosts. However, while black hosts benefited from the "smart pricing" tool, they were less likely to adopt it in the first place. As a result, the racial revenue gap continues to be substantial in the short-term rental market. Pandey and Caliskan (2021) report similar concerns in Uber's location-based pricing algorithms, where they found that more nonwhite neighborhoods in Chicago experience higher prices. If these biases persist, regulators may want to require firms to use "de-bias" pricing algorithms (see Kearns and Roth (2019) for a broader overview).

3.3 Algorithmic Collusion

In a standard reinforcement learning algorithm, there is an implicit assumption that, conditional on the demand state, the true distribution of underlying preferences is stable or stationary but unknown to the algorithm. The implementation of these algorithms for pricing (e.g., Kleinberg & Leighton, 2014; Misra et al., 2019) can be considered as dynamic field experiments that balance learning about preferences and earning profits. However, what if multiple firms are simultaneously learning each others' algorithms? An increasing body of work suggests an important distortion: Long-run prices can be above the competitive level, which has led to antitrust concerns that companies may inadvertently end up *colluding through algorithms*.

Here we emphasize that directly programming collusion is illegal and has been prosecuted by the US Department of Justice. As per a case filings, the “defendant and his co-conspirators adopted specific pricing algorithms for the sale of certain posters with the goal of coordinating changes to their respective prices and wrote computer code that instructed algorithm-based software to set prices in conformity with this agreement.”¹⁵

Having said that, we now review mechanisms that can cause independent, competing pricing algorithms to set prices that are supra-competitive.¹⁶ The extant literature tends to rely on theoretical or simulated markets, uncovering mechanisms that policy makers and regulators may want to monitor going forward. Hansen, Misra, and Pai (2021c) consider a model where firms act as local monopolists and MAB (i.e., single-state reinforcement learning) models for pricing (Section 2.3). They find that long-run prices depend on the informational value (or the signal-to-noise ratio) of the underlying pricing experiments. In markets where price experiments have low information value, the resulting long-run prices are statistically indistinguishable from Nash Equilibrium prices. However, in markets where price experiments have high information value, market prices are supra-competitive. They also find that markets with more informative experiments result in correlated experimentation across firms. Econometrically, in this setting, competitive prices are correlated unobservables in each firms' pricing algorithm, resulting in biased learning from experiments.

Calvano, Calzolari, Denicolo, and Pastorello (2020) and Klein (2021) study a setting where independent firms use Q-learning models within the class of reinforcement learning algorithms. Firms consider the current prices of all products in the market and then set the “optimal” prices in the next period. These studies show that algorithms consistently learn to charge supracompetitive prices without communicating with one another. Calvano, Calzolari, Denicolò, and Pastorello (2019) suggest that the strategies the firms learn are consistent with “stick-and-carrot” strategies, described in Abreu (1986), where firms learn that deviations from collusive prices are met with an immediate response. As a result, firms ignore short-term incentives to reduce prices and maintain a supra-competitive price.

Other mechanisms for algorithmic collusion include the following. (1) Timing: Brown and Mackay (2020) find that asynchronous responses can lead to price

dispersion and therefore non-Nash prices (above Nash in simulations); (2) Commitment: [Salcedo \(2015\)](#) show that if firms commit to a particular algorithm, it can enable supra-competitive prices; (3) Demand information: [Miklós-Thal and Tucker \(2019\)](#) and [O’connor and Wilson \(2019\)](#) study settings where competitors are unsure about the underlying demand, and they show in some conditions that information can increase the value of collusion; (4) Third-party pricing agents: [Harrington \(2020\)](#) proposes a setting where firms are unable to design their own algorithms; and if a third-party vendor designs them instead, it may design a “collusive” algorithm to maximize the two firms’ joint profit; (5) Asynchronous learning: [Asker, Fershtman, and Pakes \(2021\)](#) study the role that complexity plays in the algorithms, and they show that asynchronous learning (defined as learning limited to actions taken) can lead to monopoly prices.

Although a large and growing literature is being developed on algorithmic collusion, an important empirical question is still largely unanswered: Are these mechanisms stable in markets? Recent work in this area includes that of [Assad et al. \(2020\)](#), which studies prices in German gasoline stations before and after the stations adopted pricing algorithms. The authors compare prices in markets before and after algorithms are estimated to be adopted and find that prices increased in duopoly markets only when both competing firms are estimated to have adopted algorithms.

4. SUMMARY

In this review article, we discussed how advances in AI technology are enabling businesses to overcome barriers that pricing managers have been facing for decades. We have seen an unprecedented increase in firms’ adopting AI algorithms for pricing, both in online and offline markets. Three core areas motivate businesses to use AI pricing: (1) real-time responses to demand and supply, (2) price personalization, and (3) demand learning. In terms of their operationalization, we notice that each of these areas is characterized by a stylized empirical fact: The adoption of price algorithms tends to increase the frequency of price changes or the price dispersion across customers. Despite the benefits from a revenue management perspective, AI pricing raises timely and challenging questions about its effects on the public at large, including discrimination, fairness concerns, transparency, and market collusion between sellers.

Throughout the review, we have mentioned many areas that would benefit from further research. Here, we focus on three. One promising research area would be scaling up AI price algorithms for multi-product firms. Currently, the extant literature and the extant price algorithms deployed in the field tend to consider single-product price optimization. Imagine a firm using AI pricing algorithms to learn optimal prices for 20 products. For each product, the firm might consider 15 different prices (as in [Calvano et al., 2019](#)). In this case, a product-level analysis would consider 20 different problems with 15 prices each ($20 * 15$ in complexity), while a joint optimization problem would consider 15^{20}

different price vectors – clearly an intractable problem. Studying possible simplifications is an important step toward enabling firms to solve this problem.

A second area that is currently underrepresented in the literature is how AI pricing algorithms affect demand. Extant literature mainly considers web-scraped pricing data; meanwhile, we encourage scholars to obtain novel insights using real-time purchase (demand) data, combined with institutional knowledge about how the company actually sets prices. Consider a dataset that includes prices and demands for an online platform and that has the following features: (1) prices that vary across users and over short time intervals, and (2) a large number of potential substitute products. This new form of demand data poses new statistical challenges and requires academics to develop new econometric methods to accurately estimate demand primitives. For example, researchers can leverage knowledge about the specifics of the algorithm to control for endogenous price changes. Another challenge in this area of research is to consider the bias in AI learning algorithms when exposure to algorithms can change consumer primitives such as price elasticity (Aparicio et al., 2022). Future research can derive AI pricing algorithms that can incorporate changes in demand primitives.

A third important topic that would benefit from further research is how public policy and competition law can account for the accelerated changes brought by algorithm pricing. Research could consider the welfare implications from dynamic and personalized pricing in competitive markets; biases or demographic inputs that are unintentionally picked up or exacerbated by an algorithm; societal implications when price-setting is increasingly shifted from humans to machines (and the organizational or cultural tensions that might be created within a company); and potential tacit collusion between algorithms of rival firms. Related to these important phenomena, we note an overarching question: Should legal restrictions be placed on the implementation of algorithmic pricing? If so, when? Should policy-makers ask firms to reveal their pricing algorithms?

NOTES

1. For example, see Ma (2016); industry reports suggest that Amazon.com sells about 750 million different products <https://www.scrapehero.com/how-many-products-does-amazon-sell-march-2021> (last accessed 02/06/2022); in contrast, brick-and-mortar stores sell 20,000 different products (Anderson et al., 2015). For a broader overview see the billion prices project (Cavallo & Rigobon, 2016).

2. This section considers *cross-sectional* price discrimination, in contrast to prices that vary based on time varying demand and supply conditions or *inter-temporal* price discrimination discussed in the prior section.

3. More precisely, they show that Bayesian choice models (vs. prediction models from machine learning, such as neural networks or random forests) generate, on average, higher expected profits.

4. In this B2B setting, a prospective employer is a customer of the company [ziprecruiter.com](https://www.ziprecruiter.com).

5. More precisely, the authors estimate a Weighted Likelihood Bootstrap estimator to estimate a heterogeneous logit demand model.

6. See Huang, Ellickson, and Lovett (2021); Aparicio and Simester (2022) for empirical evidence.

7. At time t , consider all the data from time 0 to time t and estimate the parametric demand model. Consider a realization of the estimated parameters from the asymptotic distribution (frequentist statistics) or from the posterior distribution (Bayesian statistics). Based on this parameter, set the optimal prices for time t , and repeat.

8. In each round, each price is given an index or score based on two factors: (1) the observed historical profit from charging this price (earning); and (2) the number of times this price has been tried in the past (learning). The algorithm simply plays the price with the highest index (similar to Gittins, 1979), observes profit, and updates its information for the next round.

9. With discrete states and actions, rewards can be represented by a matrix (called the Q matrix), e.g., rows are states and columns are actions; and each value is the current estimate of the discounted rewards for each action in each state. The algorithm updates the relevant cell of this matrix as new data are realized. This is an “off-policy” algorithm and its implementation can consider a heuristic-based experimentation – for example, ϵ -greedy.

10. See: <https://www.nytimes.com/1999/10/28/business/variable-price-coke-machine-being-test-ed.html>.

11. See Bergen, Dutta, Guszczka, and Zbaracki (2021) for a discussion of managerial implications of these unintentional consequences.

12. See: <https://www.nytimes.com/2017/09/17/travel/price-gouging-hurricane-irma-airlines.html> and <https://www.nytimes.com/2020/03/27/us/coronavirus-price-gouging-hand-sanitizer-masks-wipes.html>.

13. See: <https://www.congress.gov/bill/117th-congress/house-bill/675>.

14. See: <https://ag.ny.gov/press-release/2020/attorney-general-james-stops-three-amazon-sellers-price-gouging-hand-sanitizer>.

15. See: <https://www.justice.gov/opa/pr/former-e-commerce-executive-charged-price-fixing-antitrust-divisions-first-online-marketplace>.

16. See Hansen, Misra, and Pai (2021b); Harrington (2018); Calvano, Calzolari, Denicolò, Harrington, and Pastorello (2020) for a more detailed summary of the literature.

REFERENCES

- Abreu, D. (1986). Extremal equilibria of oligopolistic supergames. *Journal of Economic Theory*, 39, 191–225.
- Adams, B., & Williams, K. R. (2019). Zone pricing in retail oligopoly. *American Economic Journal: Microeconomics*, 11, 124–156.
- Aghion, P., Bolton, P., Harris, C., & Jullien, B. (1991). Optimal learning by experimentation. *The Review of Economic Studies*, 58, 621–654.
- Allon, G., Cohen, M. C., & Sinchaisri, W. P. (2018). The impact of behavioral and economic drivers on gig economy workers. Available at SSRN 3274628.
- Anderson, E., Jaimovich, N., & Simester, D. (2015). Price stickiness: Empirical evidence of the menu cost channel. *The Review of Economics and Statistics*, 97, 813–826.
- Aparicio, D., Eckles, D., & Kumar, M. (2022). *Algorithmic pricing and consumer sensitivity to price variability*. Working Paper. MIT.
- Aparicio, D., Metzman, Z., & Rigobon, R. (2021). *The pricing strategies of online grocery retailers*. Working Paper No. 28639. NBER.
- Aparicio, D., & Rigobon, R. (2021). *Quantum prices*. Working Paper No.26646. NBER.
- Aparicio, D., & Simester, D. (2022). Price frictions and the success of new products. *Marketing Science*, forthcoming.
- Asker, J., Fershtman, C., & Pakes, A. (2021). *Artificial intelligence and pricing: The impact of algorithm design*. Technical Report. National Bureau of Economic Research.
- Assad, S., Clark, R., Ershov, D., & Xu, L. (2020). *Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market*. Working Paper. CESifo.
- Auer, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 397–422.

- Baker, W., Kiewell, D., & Winkler, G. (2014). Using big data to make better pricing decisions. *McKinsey and Company*.
- Bergen, M. E., Dutta, S., Guszczka, J., & Zbaracki, M. J. (2021). How AI can help companies set prices more ethically. *Harvard Business Review*.
- Borenstein, S., & Rose, N. L. (1994). Competition and price dispersion in the US airline industry. *Journal of Political Economy*, 102, 653–683.
- Brown, Z., & Mackay, A. (2020). Competition in pricing algorithms. Available at SSRN 3485024.
- Cachon, G. P., & Swinney, R. (2011). The value of fast fashion: Quick response, enhanced design, and strategic consumer behavior. *Management Science*, 57, 778–795.
- Calvano, E., Calzolari, G., Denicolò, V., Harrington, J. E., & Pastorello, S. (2020). Protecting consumers from collusive prices due to AI. *Science*, 370, 1040–1042.
- Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2019). Algorithmic pricing what implications for competition policy? *Review of Industrial Organization*, 55, 155–171.
- Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020b). Artificial intelligence, algorithmic pricing, and collusion. *The American Economic Review*, 110, 3267–3297.
- Castillo, J. C. (2020). Who benefits from surge pricing? Available at SSRN 3245533.
- Cavallo, A. (2017). Are online and offline prices similar? Evidence from large multichannel retailers. *The American Economic Review*, 107, 283–303.
- Cavallo, A. (2018). *More Amazon effects: Online competition and pricing behaviors*. Technical Report. National Bureau of Economic Research.
- Cavallo, A., Neiman, B., & Rigobon, R. (2014). Currency unions, product introductions, and the real exchange rate. *Quarterly Journal of Economics*, 129, 529–595.
- Cavallo, A., & Rigobon, R. (2016). The billion prices project: Using online prices for measurement and research. *The Journal of Economic Perspectives*, 30, 151–178.
- Cheng, Y. (2008). Dynamic pricing decision for perishable goods: A Q-learning approach. In 2008 4th International Conference on Wireless Communications, Networking and Mobile Computing (pp. 1–5). IEEE.
- Chen, L., Mislove, A., & Wilson, C. (2016). An empirical analysis of algorithmic pricing on Amazon marketplace. In Proceedings of the 25th International World Wide Web Conference (WWW 2016), Montreal, Canada.
- Chintagunta, P. K., Dubé, J.-P., & Singh, V. (2003). Balancing profitability and customer welfare in a supermarket chain. *Quantitative Marketing and Economics*, 1, 111–147.
- Cho, S., & Rust, J. (2010). The flat rental puzzle. *The Review of Economic Studies*, 77, 560–594.
- Cowgill, B., & Tucker, C. E. (2019). Economics, fairness and algorithmic bias. *The Journal of Economic Perspectives*, forthcoming.
- Cui, R., Li, M., & Zhang, S. (2022). AI and procurement. *Manufacturing & Service Operations Management*, 24, 691–706.
- Dellavigna, S., & Gentzkow, M. (2019). Uniform pricing in us retail chains. *Quarterly Journal of Economics*, 134, 2011–2084.
- Duan, J. A., & Mela, C. F. (2009). The role of spatial demand on outlet location and pricing. *Journal of Marketing Research*, 46, 260–278.
- Dubé, J.-P., & Misra, S. (2021). Personalized pricing and consumer welfare. Available at SSRN 2992257.
- Elberg, A., Gardete, P. M., Macera, R., & Noton, C. (2019). Dynamic effects of price promotions: Field evidence, consumer search, and supply-side implications. *Quantitative Marketing and Economics*, 17, 1–58.
- Ellickson, P. B., & Misra, S. (2008). Supermarket pricing strategies. *Marketing Science*, 27, 811–828.
- Feinberg, F. M., Krishna, A., & Zhang, Z. J. (2002). Do we care what others get? A behaviorist approach to targeted promotions. *Journal of Marketing Research*, 39, 277–291.
- Fisher, M., Gallino, S., & Li, J. (2018). Competition-based dynamic pricing in online retailing: A methodology validated with field experiments. *Management Science*, 64, 2496–2514.
- Gabel, S., & Timoshenko, A. (2021). Product choice with large assortments: A scalable deep-learning model. *Management Science*.
- Ganti, R., Sustik, M., Tran, Q., & Seaman, B. (2018). Thompson sampling for dynamic pricing. *arXiv preprint arXiv:1802.03050*.

- Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society: Series B*, 41, 148–177.
- Gittins, J. C., Glazebrook, K., & Weber, R. (2011). *Multi-armed bandit allocation indices* (2nd ed.). New York, NY: John Wiley & Sons.
- Hajian, S., Bonchi, F., & Castillo, C. (2016). Algorithmic bias: From discrimination discovery to fairness-aware data mining. In Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining (pp. 2125–2126). ACM.
- Hall, J., Kendrick, C., & Nosko, C. (2015). *The effects of Uber's surge pricing: A case study*. Chicago, IL: The University of Chicago Booth School of Business.
- Hansen, K. T., Misra, K., & Pai, M. M. (2021b). Collusive outcomes via pricing algorithms. *Journal of European Competition Law & Practice*, 12, 334–337.
- Hansen, K. T., Misra, K., & Pai, M. M. (2021c). Frontiers: Algorithmic collusion: Supra-competitive prices via independent algorithms. *Marketing Science*, 40(1), 1–12.
- Hansen, K., Misra, K., & Sanders, R. E. (2021a). Consumer (in) attention to expiration dates: A field study. *Available at SSRN*.
- Harrington, J. E. (2018). Developing competition law for collusion by autonomous artificial agents. *Journal of Competition Law and Economics*, 14, 331–363.
- Harrington, J. E. (2020). Third party pricing algorithms and the intensity of competition. *Available at SSRN 3723997*.
- Haws, K. L., & Bearden, W. O. (2006). Dynamic pricing and consumer fairness perceptions. *Journal of Consumer Research*, 33, 304–311.
- Hitsch, G. J., Hortaçsu, A., & Lin, X. (2021). Prices and promotions in US retail markets. *Quantitative Marketing and Economics*, 19, 289–368.
- Hoch, S. J., Kim, B.-D., Montgomery, A. L., & Rossi, P. E. (1995). Determinants of store-level price elasticity. *Journal of Marketing Research*, 32, 17–29.
- Huang, Y. (2021). Pricing frictions and platform remedies: The case of Airbnb. *Available at SSRN 3767103*.
- Huang, Y., Ellickson, P. B., & Lovett, M. J. (2021). EXPRESS: Learning to set prices. *Journal of Marketing Research*. doi:10.1177/00222437211068527
- Kahneman, D., Knetsch, J. L., & Thaler, R. (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *The American Economic Review*, 76(4), 728–741.
- Kallus, N., & Zhou, A. (2021). Fairness, welfare, and equity in personalized pricing. In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (pp. 296–314).
- Karlinksky-Shichor, Y., & Netzer, O. (2019). Automating the B2B salesperson pricing decisions: Can machines replace humans and when. *Available at SSRN 3368402*.
- Kearns, M., & Roth, A. (2019). *The ethical algorithm: The science of socially aware algorithm design*. New York, NY: Oxford University Press.
- Klein, T. (2021). Autonomous algorithmic collusion: Q-learning under sequential pricing. *The RAND Journal of Economics*, 52, 538–558.
- Kleinberg, R., & Leighton, F. T. (2014). The value of knowing a demand curve: Bounds on regret for on-line posted-price auctions. (Akamai Technologies).
- Kleinberg, J., Ludwig, J., Mullainathan, S., & Rambachan, A. (2018). Algorithmic fairness. In *Aea papers and proceedings* (Vol 108, pp. 22–27).
- Lambrecht, A., & Tucker, C. (2019). Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. *Management Science*, 65(7), 2966–2981.
- Li, Y., Gordon, B. R., & Netzer, O. (2018). An empirical study of national vs. local pricing by chain stores under competition. *Marketing Science*, 37, 812–837.
- Liu, X. (2021). *Dynamic coupon targeting using batch deep reinforcement learning: An application to LiveStream shopping*. Working Paper.
- Lu, R., Hong, S. H., & Zhang, X. (2018). A dynamic pricing demand response algorithm for smart grid: Reinforcement learning approach. *Applied Energy*, 220, 220–230.
- Ma, J. (2016). Does greater online assortment pay? An empirical study using matched online and catalog shoppers. *Journal of Retailing*, 92, 373–382.
- Mcafee, P., & Te Velde, V. (2006). Dynamic pricing in the airline industry. In T. J. Hendershott (Ed.), *Handbook on economics and information systems* (Vol. 1, Chapter 11, pp. 527–570). Amsterdam: Elsevier.

- Miklós-Thal, J., & Tucker, C. (2019). Collusion by algorithm: Does better demand prediction facilitate coordination between sellers? *Management Science*, *65*, 1552–1561.
- Misra, K., Schwartz, E. M., & Abernethy, J. (2019). Dynamic online pricing with incomplete information using multiarmed bandit experiments. *Marketing Science*, *38*, 226–252.
- Musolf, L. (2021). *Algorithmic pricing facilitates tacit collusion: Evidence from ECommerce*. Working Paper.
- O’connor, J., & Wilson, N. (2019). *Reduced demand uncertainty and the sustainability of collusion: How AI could affect competition*. Working Paper. FTC Bureau of Economics.
- Orbach, B. Y., & Einav, L. (2007). Uniform prices for differentiated goods: The case of the movie-theater industry. *International Review of Law and Economics*, *27*, 129–153.
- Pancras, J., & Sudhir, K. (2007). Optimal marketing strategies for a customer data intermediary. *Journal of Marketing Research*, *44*, 560–578.
- Pandey, A., & Caliskan, A. (2021). Disparate impact of artificial intelligence bias in ridehailing economy’s price discrimination algorithms. In Proceedings of the 2021 AAAI/ACM Conference on Ai, Ethics, and Society (pp. 822–833).
- Proserpio, D., Hauser, J. R., Liu, X., Amano, T., Burnap, A., Guo, T., & Yoganarasimhan, H. (2020). Soul and machine (learning). *Marketing Letters*, *31*, 393–404.
- Rossi, P. E., Mcculloch, R. E., & Allenby, G. M. (1996). The value of purchase history data in target marketing. *Marketing Science*, *15*, 321–340.
- Rotemberg, J. J. (2011). Fair pricing. *Journal of the European Economic Association*, *9*, 952–981.
- Salcedo, B. (2015). *Pricing algorithms and tacit collusion*. Working Paper. Pennsylvania State University.
- Seele, P., Dierksmeier, C., Hofstetter, R., & Schultz, M. D. (2021). Mapping the ethicality of algorithmic pricing: A review of dynamic and personalized pricing. *Journal of Business Ethics*, *170*, 697–719.
- Shiller, B., & Waldfogel, J. (2011). Music for a song: An empirical look at uniform pricing and its alternatives. *The Journal of Industrial Economics*, *59*, 630–660.
- Smith, A. N., Seiler, S., & Aggarwal, I. (2021). Optimal price targeting. Available at SSRN 3975957.
- Snyder, J. (2009). What’s the matter with price gouging? *Business Ethics Quarterly*, *19*, 275–293.
- Stamatopoulos, I., Bassamboo, A., & Moreno, A. (2021). The effects of menu costs on retail performance: Evidence from adoption of the electronic shelf label technology. *Management Science*, *67*, 242–256.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- The Atlantic. (2017). How online shopping makes suckers of us all. *The Atlantic*. May 2017, Jerry Useem.
- Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, *25*, 285–294.
- Van Der Rest, J.-P. I., Sears, A. M., Miao, L., & Wang, L. (2020). A note on the future of personalized pricing: Cause for concern. *Journal of Revenue and Pricing Management*, *19*, 113–118.
- Watkins, C. J., & Dayan, P. (1992). Q-learning. *Machine Learning*, *8*, 279–292.
- Weisstein, F. L., Monroe, K. B., & Kukar-Kinney, M. (2013). Effects of price framing on consumers’ perceptions of online dynamic pricing practices. *Journal of the Academy of Marketing Science*, *41*, 501–514.
- White, J. M. (2012). *Bandit algorithms for website optimization*. Sebastopol, CA: O’Reilly Media.
- ‘White House’, C. O. E. A. (2015, February). *Big data and differential pricing*.
- Zhang, D. J., Dai, H., Dong, L., Qi, F., Zhang, N., Liu, X., . . . Yang, J. (2020). The long-term and spillover effects of price promotions on retailing platforms: Evidence from a large randomized experiment on Alibaba. *Management Science*, *66*, 2589–2609.
- Zhang, J., & Krishnamurthi, L. (2004). Customizing promotions in online stores. *Marketing Science*, *23*, 561–578.
- Zhang, S., Mehta, N., Singh, P. V., & Srinivasan, K. (2021). Frontiers: Can an AI algorithm mitigate racial economic inequality? An analysis in the context of Airbnb. *Marketing Science*, forthcoming.
- Zhang, J., & Wedel, M. (2009). The effectiveness of customized promotions in online and offline stores. *Journal of Marketing Research*, *46*, 190–206.