



UvA-DARE (Digital Academic Repository)

Online Price Discrimination and EU Data Privacy Law

Zuiderveen Borgesius, F.; Poort, J.

DOI

[10.1007/s10603-017-9354-z](https://doi.org/10.1007/s10603-017-9354-z)

Publication date

2017

Document Version

Final published version

Published in

Journal of Consumer Policy

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Zuiderveen Borgesius, F., & Poort, J. (2017). Online Price Discrimination and EU Data Privacy Law. *Journal of Consumer Policy*, 40(3), 347-366. <https://doi.org/10.1007/s10603-017-9354-z>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Online Price Discrimination and EU Data Privacy Law

Frederik Zuiderveen Borgesius & Joost Poort

Journal of Consumer Policy

Consumer Issues in Law, Economics and Behavioural Sciences

ISSN 0168-7034

Volume 40

Number 3

J Consum Policy (2017) 40:347-366

DOI 10.1007/s10603-017-9354-z



Your article is published under the Creative Commons Attribution license which allows users to read, copy, distribute and make derivative works, as long as the author of the original work is cited. You may self-archive this article on your own website, an institutional repository or funder's repository and make it publicly available immediately.

Online Price Discrimination and EU Data Privacy Law

Frederik Zuiderveen Borgesius¹ · Joost Poort¹ 

Received: 30 September 2016 / Accepted: 16 June 2017 /
Published online: 15 July 2017
© The Author(s) 2017. This article is an open access publication

Abstract Online shops could offer each website customer a different price. Such personalized pricing can lead to advanced forms of price discrimination based on individual characteristics of consumers, which may be provided, obtained, or assumed. An online shop can recognize customers, for instance through cookies, and categorize them as price-sensitive or price-insensitive. Subsequently, it can charge (presumed) price-insensitive people higher prices. This paper explores personalized pricing from a legal and an economic perspective. From an economic perspective, there are valid arguments in favour of price discrimination, but its effect on total consumer welfare is ambiguous. Irrespective, many people regard personalized pricing as unfair or manipulative. The paper analyses how this dislike of personalized pricing may be linked to economic analysis and to other norms or values. Next, the paper examines whether European data protection law applies to personalized pricing. Data protection law applies if personal data are processed, and this paper argues that that is generally the case when prices are personalized. Data protection law requires companies to be transparent about the purpose of personal data processing, which implies that they must inform customers if they personalize prices. Subsequently, consumers have to give consent. If enforced, data protection law could thereby play a significant role in mitigating any adverse effects of personalized pricing. It could help to unearth how prevalent personalized pricing is and how people respond to transparency about it.

Keywords Price discrimination · Cookie · Data protection law · General Data Protection Regulation · Behavioural targeting · Personalized communication

The original version of this article was revised: The last name of “Frederik Zuiderveen Borgesius” should have been “Zuiderveen Borgesius”.

✉ Joost Poort
poort@uva.nl

Frederik Zuiderveen Borgesius
f.j.zuiderveenborgesius@uva.nl

¹ IViR Institute for Information Law, University of Amsterdam, P.O. Box 1030, 1000 BA Amsterdam, The Netherlands

On the internet, companies are able to personalize prices based on information about consumers. For instance, a company could categorize consumers as “high spenders” if their browsing profile or purchase history suggest that they often buy expensive goods. Already before the turn of the millennium, a book on database marketing explained that firms should not treat all customers the same: “The buyer-seller relationship is not a democracy. All customers are not created equal. All customers are not entitled to the same inalienable rights, privileges, and benefits. (...) That means some customers must earn ‘better treatment’ than others, whatever that means. If you cannot accept this undemocratic fact, quit reading and close the book, right now. Database relationship marketing is not for you.” (Newell 1997, p. 136).

This paper explores personalized online price discrimination (“personalized pricing”) from both a legal and an economic perspective. The aim is neither to contribute to economic scholarship of price discrimination itself nor to deal with all areas of law that may be relevant to personalized pricing. Rather, the aim is to assess how personalized pricing works in practice, what can be learnt about it from economics, and how this can be linked to people’s general dislike of personalized pricing. Given the fact that personalized pricing is shown to raise concerns from the perspective of social welfare and its distribution and is met with much scepticism by society, a valid question is if existing law can be of help. This paper discusses whether EU data privacy law applies to online price discrimination and if so, what implications that has.

The next section sketches the current practice of personalised pricing and outlines the sources of the data to personalise prices on. Subsequently, it delineates the scope of personalised pricing for the purpose of this paper. Next, we discuss the basic economics of price discrimination, in particular in an online environment. We also show that opinions of non-economic scholars and consumers differ from those of economists on the question of whether personalised pricing is desirable. We analyse how a widespread dislike of personalised pricing can be understood from an economic perspective.

Then, we turn to data protection law. We argue that personalised pricing generally entails the processing of personal data. Data protection law requires a company to inform people about the purpose of processing their personal data. Hence, companies must say so if they personalise prices; such transparency may impact the practice and outcome of personalised pricing. Moreover, companies must generally obtain the customer’s informed consent before they can legally use personalised pricing. Finally, we mention some caveats about the usefulness of data protection law in the context of personalised pricing. We also give suggestions for further research, and provide concluding thoughts.

Personalized Pricing

Online price discrimination or personalized pricing can be described as *differentiating the online price for identical products or services partly based on information a company has about a potential customer*.¹ Among the earliest scholars to write about online price discrimination are Odlyzko (1996, 2003) and Baker et al. (2001). The latter suggest that smart sellers use online price discrimination: “Just as it’s easy for

¹ A more rigorous economic discussion of price discrimination will be provided in “Price Discrimination: Economics” section.

customers to compare prices on the Internet, so is it easy for companies to track customers' behaviour and adjust prices accordingly." They add that "[t]he Internet also allows companies to identify customers who are happy to pay a premium." Thus, an online shop could charge higher prices to high-spending or price-insensitive people and vice versa.

In 2000, it was widely reported that Amazon offered different visitors different prices (e.g., BBC News 2000). When a regular customer deleted his computer's cookies, he saw the price of a DVD drop. Hence, it appeared that customers who previously ordered from Amazon were shown a higher price for a product than new customers. Many people reacted angrily and raised the issue of fairness (e.g., Krugman 2000). Amazon hastily issued a press release stating that it was merely experimenting with random discounts and gave a refund to people who paid a price above the average. Amazon's CEO Jeff Bezos said: "We have never tested and we never will test prices based on customer demographics" (Amazon 2000).

A second example concerns office supplies sold by Staples.com. Valentino-Devries et al. (2012) showed that Staples charged people in different areas different prices, based on their IP address. This had the effect, likely unintentional, that people from high-income areas paid less. Similarly, Mikians et al. (2013) found several examples of online shops that charge customers from different regions different prices, while Hannak et al. (2014) found examples of companies that offer discounts to mobile users and to people who are logged in to sites.

New data analysis technologies, often summarized as "big data," give new possibilities for tailored pricing. A White House report notes: "Many companies already use big data for targeted marketing, and some are experimenting with personalized pricing, though examples of personalized pricing remain fairly limited." The report adds that "[t]he increased availability of behavioural data has also encouraged a shift from (...) price discrimination based on broad demographic categories towards personalized pricing." (Executive Office of the President of the United States 2015, p. 2–4).

Notwithstanding the examples above, personalized pricing seems to be relatively rare. As Narayanan (2013) puts it: "The mystery about online price discrimination is why so *little* of it seems to be happening." One reason may be that companies are hesitant to personalize prices because they fear consumer backlash as many people find price discrimination unfair (Office of Fair Trading 2010, p. 48; Executive Office of the President of the United States 2015, p. 13). Odlyzko (2009, p. 50) writes: "The main constraint on price discrimination comes from society's dislike of the practice." He adds that it is unclear "what forms of price discrimination society will accept. So we should expect experimentation, hidden as much as sellers can manage, but occasionally erupting in protests, and those protests leading to sellers pulling back, at least partially. And occasionally we should expect government action, when the protests grow severe."

As Odlyzko suggests, personalized pricing may happen covertly. It is difficult to detect whether companies actually adjust prices based on an individual's characteristics. Prices may fluctuate for generic reasons. For instance, a shop may try out different prices at different times, like Amazon claimed to be doing in 2000. And, airline ticket prices fluctuate based on many factors such as the days before departure and the seats left. Hence, if you delete your cookies and see a different price for an airline ticket, it is not proven that the airline adapts prices to people's cookie profiles.

You might see a price difference that is unrelated to your cookie (Vissers et al. 2014), as Amazon claimed was the case in the example above.

When a company *does* personalize prices online, it can do this on the basis of information from a wide variety of sources, containing different sorts of data:

- Information can be *voluntarily and knowingly provided by a customer*, e.g., when he or she created a customer account for a previous purchase. Such a customer can be recognized when he or she logs in. Information provided by customers can be any kind of personal and commercially interesting information, such as name, address, date of birth, type of credit card, previous purchases, wish lists, or gender.
- Information can also be *involuntarily and unknowingly obtained from a customer or website visitor*. For instance, an online shop can recognize a consumer on the basis of her IP address or a cookie with a unique identifier. An IP address can also be used to derive the country and region where a consumer lives and the type of internet provider she has. The type and speed of the internet connection (mobile or fixed, cable, copper or fibre) and the browser or computer type can also be acquired.
- A third broad category is information *obtained from third parties* such as advertising networks or gathered by observing a customer's online behaviour through *affiliated ad networks* or on *affiliated websites*. Using cookies or other technologies, an ad network can recognize Internet users when they visit websites and can build a profile of them. When they visit an online shop, the shop could adapt prices based on their profile.

All these sorts of information, in particular the latter two, can be used to make inferences about income, age, and social strata, which may be informative about a website visitor's willingness to pay. For instance, someone living in a rich neighbourhood is more likely to be rich. Likewise, past purchases and browsing behaviour can be used to make assumptions about someone's expected willingness to pay for a product. Note that such assumptions or inferences need not be fully correct for them to work: Poor or stingy people may live in rich neighbourhoods, and many more people browse the web for Ferraris than can actually afford them. Nevertheless, pricing algorithms based on such data may still increase a website's revenues. In sum, it is as yet unclear how common personalized pricing is and more empirical research to assess this is needed. But, the technology to personalize prices is there, and if companies can use it legally to raise their profits, they can be expected to do so.

Price Discrimination: Economics

Definition and Types of Price Discrimination

In the last section, online price discrimination was defined as differentiating the online price for identical products or services partly based on information a company has about a potential customer. Simple examples, which many people feel comfortable with, are a reduced conference fee for academics or doctoral students as opposed to participants from commercial entities, or a reduced fee for senior citizens or children when booking theatre tickets online. However, this definition fails to account for differences in *cost* which serving different customers may entail and which may justify price differences. Transportation or delivery costs can differ between customers, for instance, and including such costs in the total charge would

generally not be considered price discrimination. Similarly, differences in legal regimes can translate into diverse cost levels for the same service between different EU member states. Other more personal characteristics may also give rise to cost differences resulting in price differences, for instance in insurance and credit markets. Based on demographic data or a person's track record, he or she may have a higher probability to cause a traffic accident, fall ill, become unemployed, or default on a loan. By consequence, the cost of providing insurance or credit will differ. These cost differences will generally justify price differences that most authors would not consider price discrimination.²

An economically more proper definition of price discrimination by Stigler is "the sale of two or more similar goods at prices that are in different ratios to marginal costs" (Stigler 1987, p. 210). Indeed, under this definition, the examples above of price differences that purely stem from cost differences would not qualify as price discrimination.

Stigler's definition of price discrimination can be criticized for being vague by using the word "similar." However, an advantage of this vagueness is that the definition also covers versioning of products. A classic example of versioning is selling a book in a paperback and a hardcover edition. Publishers usually make more profit on the hardcover. A notorious example of versioning concerns IBM, which added a microchip to its popular professional LaserPrinter to reduce the printing speed and subsequently sell it for half the price on the consumer market (e.g., Preston McAfee 2008). Other examples are versions of the same software that have slightly different features, cars with different engine powers, etc. All these product sets can be considered similar, and in such examples, the version that has superior features or quality is typically sold with a higher profit margin.

For price discrimination to work, three conditions need to be satisfied: A seller must be able to distinguish between customers to know which price to charge to whom, he must have enough market power to be able to set prices above marginal costs, and resale must be impractical, costly, or forbidden to prevent arbitrage between customers (e.g., Varian 1989). For online sales, these conditions are often met: Distinguishing customers is possible with great accuracy. Various internet sellers have very high market shares in their relevant market which are likely to give them at least some market power, while market power may also derive from switching costs or lack of transparency in the market. And, resale is often impossible (in the case of airline tickets or hotel rooms for instance) or relatively costly. Combined with the ease of adapting prices unnoticed, this suggests that a lot of price discrimination should occur online.

A classic distinction is between first-, second-, and third-degree price discrimination (Pigou 1932). *First-degree price discrimination* refers to a situation in which each consumer is charged an individual price equal to his or her maximum willingness to pay. For first-degree price discrimination, the seller needs precise information about the buyer's willingness to pay (the reservation price). First-degree price discrimination enables the sellers to extract all consumer surpluses. In practice, such an extreme form of price discrimination will never occur, as sellers cannot learn the buyer's exact reservation price. First-degree price discrimination serves as a stylized benchmark to evaluate other pricing schemes.³

Second-degree price discrimination refers to pricing schemes in which the price of a good or service does not depend on characteristics of the customer but on the quantity bought. Such

² Not translating such cost differences into price differences can cause insurance markets to collapse due to adverse selection: Only bad risks will buy insurance at uniform prices.

³ Sometimes, individual prices that do not equal each individual's willingness to pay are also referred to as first degree price discrimination. Particularly in the context of online personalized pricing, this use of the term is confusing and hence avoided in this paper.

schemes are also called “non-linear pricing” and may involve a quantity discount, or a two-part tariff with a fixed fee and a variable fee. For example, in the cinema, popcorn is often cheaper (per gram) if you buy a larger box. For second-degree price discrimination, the seller does not need information about the buyer, as buyers self-select: They choose a different price by choosing a different quantity.

Certain types of loyalty schemes are sometimes also characterized as second-degree price discrimination, namely when a customer gains credits or discounts with past purchases. If a loyalty scheme only amounts to a quantity discount over time, it can indeed be seen as a type of second-degree price discrimination. However, loyalty schemes are also used for other sales strategies such as cross-selling or up-selling, i.e., to sell additional or more profitable products and services to existing customers. As a loyalty scheme is an excellent tool for making customer profiles, it can also be used for personalized pricing. This implies that loyalty schemes should also be considered within the broad scope of third-degree price discrimination.

In *third-degree price discrimination*, prices (or more properly price-to-marginal-cost-ratios) differ between groups or types of buyers. This type of price discrimination is widely used and often uncontroversial: Many companies offer discounts to children, students, or elderly. A reduced conference fee for academics is an example of third-degree price discrimination too. A company could also charge people from different geographical areas different prices. For instance, medicines or college textbooks could be sold at lower prices in developing countries.

For third-degree price discrimination, it is not necessary to recognize individual buyers: Sellers only need to know the characteristic of the buyer that is used to discriminate prices. However, to distinguish types of buyers, sellers often use unique identifiers such as a student card with a student number and photo or even a formal ID card. Uniquely identifying customers helps to satisfy two of the key conditions for price discrimination to work: distinguishing between buyers and preventing ineligible customers to obtain a discount by arbitrage.⁴ Nevertheless, such practices often lead to over-identification: More precise identification takes place than third-degree price discrimination requires.

Online price discrimination will typically work similarly: A cookie, IP address, or user log-in information enables an online company to identify a customer.⁵ Like the student ID, this unique identification will generally not be the purpose but a means to the end of third-degree price discrimination, e.g., to distinguish between broader categories, for instance high and low spenders. However, compared to selecting students on the basis of a student ID card, an online profile can be much more detailed and can allow for much more refined price discrimination. By doing so, online third-degree price discrimination can, at least in theory, be pushed towards a seller’s holy grail of perfect or first-degree price discrimination, under which all consumer surplus is extracted to the benefit of the seller.

Versioning is outside the scope of this paper. If sellers cannot easily identify groups of customers for price discrimination, they can tempt customers to self-select by using *versioning*.⁶ For example, customers with a higher willingness to pay could be tempted to buy hardcover books instead of paperbacks. Airlines aim to separate price-insensitive business travellers from price-sensitive leisure travellers by requiring a stay-over on Saturday night for cheaper tickets. Versioning could also be used to offer customers with a certain profile only the more expensive

⁴ For example, a photograph on a student ID prevents other people from using the ID card.

⁵ There may be some noise due to different people using the same computer or Internet connection, or conversely the same person using several devices, buying a new device, or deleting her cookies.

⁶ For some striking examples of gender-based versioning, see De Blasio and Menin (2015).




Fig. 1 Stylized example of monopolist applying price discrimination

versions of a product. Such practices generally do not require the seller to have specific information about the buyer.

Price steering also falls outside the scope of this paper, as it does not involve price differentials between customers. Price steering is “personalising search results to place more or less expensive products at the top of the list” (Hannak et al. 2014). Hence, a website suggests a certain product and price to a consumer, but the consumer can still buy another product. For instance, the travel site Orbitz showed Apple users more expensive hotels than it showed to PC users (Mattoli 2012). Note, however, that since the use of specific information about potential buyers will also be required for price steering, conclusions about the applicability of data protection law may transfer to it.

Welfare Effects of Price Discrimination

From a welfare economic perspective, there can be good arguments in favour of price discrimination, as it can benefit both buyers and sellers. To illustrate this, imagine a monopolist who supplies a good with constant marginal production cost of €5 to a market of 100 customers (Fig. 1). Half of these customers, the high spenders, are willing to pay an amount between €15 and €20 for this product; the low spenders between €8 and €10.⁷ Under uniform pricing, the monopolist would consider two price points. At P_{low} = €8, all customers are buying. The monopolist sells 100 units (Q_{low} = 100), making a profit of €3 per unit, €300 in total. Low spenders have a consumer surplus (the difference between their willingness to pay and what they pay) of between €0 and €2, and high spenders between €7 and €12. Total consumer surplus at this price is 50 × (€9–8) + 50 × (€17.5–8) = €525. Total welfare, the sum of consumer surplus and profits, equals €825.

⁷ For simplicity, both groups are assumed to be homogeneously distributed within this range. That is, the average willingness to pay within the first group is 17.5 € (€ 15/2 + 20/2) and within the second group it is 9 € (€ 8/2 + 10/2).

At a price $P_{\text{high}} = €15$, the monopolist sells only 50 units to high spenders but makes a profit of €10 per unit, yielding a profit of €500. Hence, without price discrimination, the monopolist will set a price of €15, sell 50 units and leave low spenders unsupplied, even though they are willing to pay more than the marginal costs of production. Consumer surplus in this scenario is $50 \times (\€0 + 5)/2 = €125$. Total welfare equals €625, which is less than in the previous scenario.

What will happen if the monopolist can price discriminate and charge €15 to high spenders and €8 to low spenders? Total profits become $50 \times €10 + 50 \times €3 = €650$, making the monopolist better off. Consumer surplus becomes $50 \times (\€17.5 - 15) + 50 \times (\€9 - 8) = €175$. Total welfare becomes €825, leaving both the monopolist and the consumers better off in comparison to the uniform price of €15 that is optimal for the monopolist.

Total welfare under price discrimination (€825) is equal to total welfare at the lower uniform price of €8. This is no coincidence. In a simple case such as this, welfare is created by the transactions that take place in the market, while the price paid only affects the distribution of welfare between the consumers and the producer. Therefore, any price between €5 and €8 generates the same (maximum) amount of total welfare. When a non-discriminating price is raised above €8, welfare is destroyed because customers who have a willingness to pay above marginal costs are not served by the market which entails so-called dead-weight-losses.⁸

The example illustrates that price discrimination can reduce deadweight losses that result from consumer heterogeneity: Uniform prices often leave some consumers unserved, even though they have a willingness to pay which exceeds marginal costs. This causes a loss of welfare both for consumers and sellers, relative to a situation in which it would be possible to charge these consumers a price that is only slightly above marginal costs. The potential for price discrimination is particularly high in markets with high fixed costs and low marginal costs: Price discrimination can help the seller to regain its fixed costs without (substantial) dead-weight-losses.

On the other hand, for some customers, price discrimination will lead to higher prices than a uniform price. Hence, price discrimination deprives some consumer groups of some of their consumer surplus. The more refined the price discrimination scheme that the seller uses, the more he can deprive consumers of their consumer surplus. For instance, one could envisage the monopolist in Figure 1 discriminating further within the group of high spenders, to extract more surplus. This is not so much a concern for the aggregate welfare created in a market, but it is for the distribution of the welfare between producers and consumers.

Moreover, price discrimination can enhance market power: In the example above, the uniform price that was optimal for the monopolist would leave half the market unsupplied, thereby leaving room for competitors to serve the low end of the market. By using price discrimination, the monopolist can serve the entire market, with the result that it can increase its economies of scale and network effects. Thus, price discrimination can help to monopolize a market and to make market entry unattractive for competitors.

There is a large literature on the outcomes and welfare effects of price discrimination in various different competitive settings and under various assumptions about consumer demand, information that consumers and producers have, producer's ability to commit to prices, etc. For an overview, see for instance Varian (1989) and Armstrong (2006). These welfare effects turn out to be ambiguous. Generally, for price discrimination to be welfare enhancing, it must lead to a substantial increase in total output by serving markets that were previously unserved. For instance, in the stylized case above, price discrimination can lead to a welfare improvement for both

⁸ A price below 5 € could attract new customers with a willingness to pay below the production cost. Serving these customers destroys welfare.

consumers and sellers. But even then, price discrimination will lead to a loss of welfare for some consumers.

When price discrimination does not lead to substantial market expansion, it often reduces total consumer surplus to the benefit of producer surplus. Price discrimination may even lead to a net welfare loss, when producers gain but consumers lose more. And sometimes, even sellers can suffer a welfare loss, due to intensified competition. The closer personalized pricing approaches first-degree price discrimination, the more it will extract welfare away from consumers and towards producers. Finally, price discrimination often introduces transaction costs for consumers and producers.

In conclusion, a welfare economic stance towards price discrimination is ambiguous. Under the right circumstances, price discrimination can increase total welfare and can even be averagely beneficial for consumers, provided it leads to a substantial increase in total output. On the other hand, price discrimination can lead to a transfer of welfare from consumers to sellers or even to a reduction of total welfare. In any case, consumers with a high willingness to pay will most probably be worse off under price discrimination.

People's Attitude Towards Personalized Pricing

The welfare economic analysis in the previous section suggests that a case-by-case analysis is needed since price discrimination can be either beneficial or detrimental to individual consumers and total consumer welfare, depending on the circumstances. However, many people regard personalized pricing generally as unfair or manipulative. In a US survey, Turow et al. (2005), p. 4) “found that [American adults] overwhelmingly object to most forms of behavioural targeting and all forms of price discrimination as ethically wrong.” In another US survey, 78% of the respondents did not want “discounts (...) tailored for you based on following (...) what you did on *other* websites you have visited” (Turow et al. 2009, p. 15). Hence, many people seem to dislike online price discrimination, even if it leads to receiving discounts.

People may dislike personalized pricing for various reasons. Through the lens of behavioural economics, such dislike could be associated with the concept of loss or regret aversion (e.g., Loomes and Sugden 1982). People avoid situations that could lead to regret. Hence, people probably object to a situation in which they would have been offered a better price if they had used a different browser or computer, or deleted their cookies.

It appears that many people are uncomfortable with personalized pricing because it can happen surreptitiously (Executive Office of the President of the United States 2015). Consumers are subject to information asymmetries with respect to pricing, as they do not know in which price category they are placed, or how to influence that. In this respect, online price discrimination differs from, say, a signposted discount for the elderly. After finding out about online personalized pricing, a consumer could perceive a loss or regret for not having circumvented it. Although there would only be cause for regret when personalization leads to higher prices, loss aversion implies that people are inclined to give this scenario more weight. The mere fear or suspicion of paying a premium could cause people to dislike personalized pricing.

Turow (2011) warns that modern data analysis technologies enable a company to classify people as “targets” or “waste,” and treat them accordingly. A company could offer discounts to people classified as high spenders to lure them to become regular customers. And, it could refrain from offering discounts to low-spending or price-sensitive people, if it expects that they will not become valuable customers anyway.

There can also be more generic reasons for people's dislike of personalized pricing. To enable personalized pricing, companies might collect more information about internet users (Turow 2011, p. 108–110). Odlyzko (2009) suggests that “the growing erosion of privacy can be explained by companies preparing to exploit the growing opportunities for price discrimination.” Many questions are still open regarding people's aversion to personalized pricing. Is the aversion partly triggered by the word discrimination? Do people dislike personalized discounts as much as they dislike personalised premiums? Do people accept price discrimination on some grounds or for some purposes, based on some underlying notion of fairness, and not for the mere purpose of raising profits? Do people just need time to get used to this new kind of price discrimination?

The distributional aspect of price discrimination, favouring companies over consumers, may also worsen the public opinion towards it (Strandburg 2013, p. 135). As Miller (2013, p. 69) points out: “Even the most hardheaded economist ought to concede that practices that increase overall social welfare but harm most consumers raise serious ethical concerns.” And Zarsky (2004, p. 53) notes: “[v]endors can overcharge when the data collected indicates that the buyer is indifferent, uninformed or in a hurry.”

From a macro-economic perspective, it would be bad for the economy if personalized pricing causes consumers to lose trust in online sellers in general (Office of Fair Trading 2010). Moreover, a mere suspicion of personalized pricing could increase search costs, lead to inefficient switching (Miettinen and Stenbacka 2015), and burden ordinary everyday transactions with a need to shop around to make sure they get the best deal.

Data Protection Law Usually Applies to Personalized Pricing

Data protection law is a legal tool that aims to ensure that the processing of personal data happens fairly, lawfully, and transparently (Art. 5, 1 (a) GDPR; Art. 8, Charter of Fundamental Rights of the European Union). We argue that data protection law generally applies to personalized pricing. Data protection law requires transparency regarding personalized pricing and generally requires companies to obtain the consumer's consent for such pricing.

Data protection law grants rights to people whose data are being processed (data subjects), and imposes obligations on parties that process personal data (data controllers, limited to companies in this paper). In this paper, we mostly refer to the General Data Protection Regulation, or GDPR (EU 2016/679), which will replace the 1995 Data Protection Directive in May 2018. Below, we discuss the personal data definition, the legal basis for personal data processing (“[Data Protection Law and a Legal Basis for Processing](#)”), and the rules on profiling and automated decisions (“[Data Protection Law and Automated Decisions](#)”).

European data protection law is triggered when “personal data” are processed. Almost anything that can be done with personal data, such as storing and analysing data, falls within the definition of “processing” (Art. 4(2) GDPR). Personal data are “any information relating to an identified or identifiable natural person (‘data subject’)” (Art. 4(1) GDPR). Typical examples of personal data are names, personal email addresses, and identification numbers.

The following hypotheticals illustrate that data protection applies to many types of personalized pricing.⁹ Suppose Nile.com is an online shop that sells consumer goods. When customers register for the site, they give Nile.com their name, personal email address, and delivery address. This is the

⁹ For the hypotheticals, we draw inspiration from Zarsky (2002).

type of information described in “[Personalized Pricing](#)” section as being voluntarily and knowingly provided by a customer.

Alice often buys at Nile.com. She only buys expensive luxury goods, never searches for cheaper alternatives on the site, and never looks at second-hand offerings. Nile.com correctly concludes that Alice is price-insensitive. When Alice logs in again, Nile.com increases all prices that Alice sees with 10%—without informing her. Alice does not realize she pays a premium and continues to buy luxury goods at Nile.com. The 10% extra that Alice pays is pure profit for Nile.com.

Bob is also registered as a Nile.com customer, but he rarely buys there. When Bob visits the site he always spends hours comparing different products and searching for second-hand offerings. Nile.com correctly concludes that Bob is very price-sensitive. Unbeknownst to Bob, Nile.com decides to offer him a 10% discount on all prices. The personalized discount leads Bob to buy more products at Nile.com. Per product sold to Bob, Nile.com makes less profit. But, since Bob now buys more products, in the aggregate, the price discount still translates into more profit for Nile.com, much like in the example in “[Price Discrimination: Economics](#)” section.

As Nile.com recognizes Alice and Bob on the basis of their log-in information (name, email address, etc.), Nile.com processes personal data. Hence, data protection law applies to the price discrimination examples regarding Alice and Bob. Say that Carol is a new customer for Nile.com, who has not registered for the site. Nile.com operates many websites and can follow a visitor to those websites through a cookie with a unique identifier (Zuiderveen Borgesius 2015, p. 15–52). This situation refers to the second and third way of acquiring information about a customer mentioned in “[Personalized Pricing](#)” section, which may lead to just as unique, specific, and effective profiles but perhaps without traditional personally identifying information such as name and address. Nile.com may not know Carol’s name, but recognizes her as the person behind the cookie with ID xyz.

By observing Carol’s browsing behaviour, Nile.com learns a lot about her. Nile.com knows that the person behind cookie xyz often visits price comparison websites and buys mostly cheap or second-hand products. Carol also visited a website with information on debt relief. Nile.com concludes that the person behind cookie xyz is price-sensitive and in some financial trouble. Therefore, when the person behind cookie xyz visits the Nile.com shopping website, Nile.com shows that person prices with a 10% discount but disables the option to buy now and pay later.

Inferences about consumers may also be incorrect. Nile.com also tracks Dave’s browsing behaviour with a tracking cookie. Dave is a community worker visiting websites on debt relief professionally, without being in financial problems himself. But, Nile.com concludes that Dave has financial problems. Nile.com disables the buy-now-pay-later option for him. Incorrect inferences could also lead to higher prices for price-sensitive and less well-to-do buyers. For instance, an online web shop could see, on the basis of a customer’s IP address, that he or she lives in a wealthy neighbourhood. But, the customer could be a poor student renting a room in that wealthy neighbourhood.

Do the Carol and Dave examples concern personal data processing? European Data Protection Authorities say that a cookie with a unique identifier tied to an individual qualifies as personal data (Article 29 Working Party 2010). This is because such cookies “enable data subjects to be ‘singled out’, even if their real names are not known” (Article 29 Working Party 2010). In line with various other scholars (e.g., De Hert and Gutwirth 2008), we agree that cookies with unique identifiers should be seen as personal data (see also Zuiderveen Borgesius 2016).

It seems likely that the Court of Justice of the European Union would also agree with Data Protection Authorities. The Court favours a broad interpretation of the concept of personal data. For instance, in 2016, the Court held that a dynamic IP address of a website visitor can be a piece of personal data for a website publisher under certain circumstances (CJEU 2016, Breyer).

Moreover, the GDPR explicitly says that “an online identifier” can identify a person: “personal data’ means any information relating to an identified or identifiable natural person (‘data subject’); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an *online identifier* or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person” (Art. 4(1), GDPR). The GDPR’s preamble confirms that singling out can be a means to identify an individual (recital 26).

One could think of a, somewhat far-fetched, hypothetical in which parts of a personalized pricing process fall outside the scope of the GDPR. Suppose that an online shop charges 10% extra to all website visitors who use an iPhone. Adapting the price on the basis of the consumer’s phone may, by itself, not constitute personal data processing. However, the hypothetical has limited practical value. Suppose Erin has an iPhone and visits the online shop. As soon as the shop ties the 10% extra price (perhaps not personal data) to Erin’s customer profile (personal data), the shop processes personal data.

In conclusion, the Alice and Bob examples clearly concern personal data processing. We argue that the Carol and Dave examples (uniquely identifying cookies) also concern personal data processing, in line with the GDPR and most likely also in line with the CJEU. And, even the example of Erin will likely concern personal data processing. Hence, EU data protection law applies to most if not all types of personalized pricing, because most personalized pricing entails the processing of personal data. The fact that data protection law applies does not imply that processing personal data is prohibited. But, if a company processes personal data, it must comply with the data protection rules. For instance, the company may only process personal data fairly, lawfully, and transparently. The transparency requirements are discussed in the next section.

Data Protection Law and Transparency

One of data protection law’s main tools to foster fairness is the requirement that data processing happens transparently. The GDPR lists the “information to be provided” by the company in detail. A company must, for instance, provide information regarding its identity and “the purposes of the processing” and must provide more information when necessary to guarantee fair processing (Art. 13 and 14 of the GDPR). Hence, a company must inform customers (data subjects) if it processes personal data with the purpose of personalising prices. As the examples regarding Alice and Bob concern personal data processing, Nile.com must inform them about personalising prices.

Compared to the 1995 Data Protection Directive, the GDPR contains more detailed transparency requirements and requires, for instance, information “in a concise, transparent, intelligible and easily accessible form, using clear and plain language” (Art. 12).

Even if, contrary to what European Data Protection Authorities and many scholars say, the Carol and Dave examples did not concern personal data processing, Nile.com would still have to inform Carol and Dave about the processing purpose. The ePrivacy Directive (Art. 5, 3) only allows storing a cookie on somebody’s computer, after that person has given consent for the cookie, “having been provided with clear and comprehensive information, in accordance with

[general data protection law], inter alia, about the purposes of the processing.” That informed consent requirement applies to many types of online identifiers, such as cookies. Hence, the ePrivacy Directive requires Nile.com to obtain Carol’s informed consent before it stores a cookie on her computer. If one of the purposes of processing the cookie is personalising prices, Nile.com must say so. The requirement to explain the purpose of a cookie applies, regardless of whether personal data are processed. A proposal for an ePrivacy Regulation, which should replace the current ePrivacy Directive, was published early 2017 (European Commission 2017). Under the proposed ePrivacy Regulation (Art. 8), a similar conclusion would be reached as under the current ePrivacy Directive.

It seems plausible that a company would prefer not telling customers about personalized prices, especially when they pay a higher price. The survey material mentioned in “[People’s Attitude Towards Personalized Pricing](#)” section indicates that if an online shop tells people it personalizes prices, people will react negatively and might therefore search for better prices or more privacy elsewhere. Currently, there seems to be no online shops that inform their customers that they personalize prices, even though the 1995 Data Protection Directive (Art. 10 and 11) requires, like the GDPR, transparency. Several explanations are possible. One possibility is that companies that are subject to EU data protection law never personalize prices. Another possibility is that some of these companies do personalize prices but do not comply with data protection law’s transparency requirements. Or, a company might think that saying, for instance, “we use personal data to offer our customers better personalised services,” is a sufficient disclosure if the company personalises prices.

However, such a broad description of the processing purpose would not comply with data protection law. Data protection law requires that companies formulate and disclose a “specified” and “explicit” purpose (Art. 5, 1(b) GDPR; Art. 5, 1(b) Data Protection Directive). European Data Protection Authorities emphasize that “Information must describe the purposes and the categories of data processed in a clear and accurate manner” (Article 29 Working Party 2012, p. 5). And “a purpose that is vague or general, such as for instance ‘improving users’ experience,’ ‘marketing purposes,’ ‘IT-security purposes’ or ‘future research’ will – without more detail – usually not meet the criteria of being ‘specific.’” (Article 29 Working Party 2013, p. 16).

In sum, the transparency requirement in data protection law requires companies to inform customers about personalising prices. If a main reason for people to dislike price discrimination is that it happens surreptitiously, mandated transparency does seem a logical policy answer. Transparency about which companies engage in price personalization could mitigate this information asymmetry: Consumers may choose online shops that do not personalize prices or delete their cookies if that gives them a better deal.

Data Protection Law and a Legal Basis for Processing

Another core principle of data protection law is the requirement that companies that process personal data have a legal basis for doing so. The GDPR exhaustively lists six possible legal bases for personal data processing; three of those legal bases could be relevant in the private sector and are discussed here. Even if a company has a legal basis for processing, it must comply with all other requirements of the GDPR.

First, a company can have a legal basis for personal data processing if “the data subject has given consent to the processing of his or her personal data for one or more specific purposes” (Art. 6(a), GDPR). However, given the dislike a majority of the people have against personalised

pricing (see “People’s Attitude Towards Personalised Pricing” section), it does not seem plausible that many would consent to this. A company should not assume that people consent if they fail to object. The requirements for valid consent are strict. For instance, the GDPR requires a “freely given, specific, informed and unambiguous indication of the data subject’s wishes by which he or she, by a statement or by a clear affirmative action, signifies agreement to the processing of personal data relating to him or her” (Art. 4(11); see also Art. 7).

Second, a legal basis could follow from the *necessity for contract performance* (Art. 6(b) GDPR). A company may process personal data if “processing is necessary for the performance of a contract to which the data subject is party or in order to take steps at the request of the data subject prior to entering into a contract.” For example, if a consumer buys a book online, the consumer’s address (a piece of personal data) may be necessary to deliver the book. And for such transactions, it is often necessary to process credit card details (Article 29 Working Party 2014, WP217, p. 16). But, it is unlikely that personal data processing for personalised pricing could be based on this legal basis. Data Protection Authorities tend to regard “necessity” as a high hurdle. The fact that a company sees personal data processing as useful or profitable does not make the processing “necessary” (Kuner 2007, p. 234–235; Article 29 Working Party 2014, WP217, p. 16 ECtHR 1983). The Court of Justice of the European Union also interprets necessity restrictively (CJEU 2008, Huber).

Third, a legal basis may lie in the balancing provision (Art. 6(f), GDPR). This allows personal data processing if it is “necessary for the purposes of the legitimate interests pursued by the controller or by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data.” In brief, a fair balance must be struck between the company’s interests (profit, for instance) and the consumer’s interests. Again, it is unlikely that Data Protection Authorities would see processing for personalised pricing as *necessary*. Moreover, they would probably say that the consumer’s interests override the company’s interests. Indeed, the Article 29 Working Party gives the following example of a practice that cannot be based on the balancing provision: “Lack of transparency about the logic of the company’s data processing that may have led to de facto price discrimination based on the location where an order is placed, and the significant potential financial impact on the customers ultimately tip the balance even in the relatively innocent context of take-away foods and grocery shopping” (Article 29 Working Party 2014, WP217, p. 32). Hence, generally, the only legal basis for personal data processing is the data subject’s consent (Steppe (2017) arrives at a similar conclusion).

For special categories of personal data (sometimes called “sensitive data”), the GDPR contains stricter rules. Special categories of data are “personal data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, and the processing of genetic data, biometric data for the purpose of uniquely identifying a natural person, data concerning health or data concerning a natural person’s sex life or sexual orientation” (Art. 9(1) GDPR). In principle, processing such special categories of data is prohibited. There are exceptions, for instance for health care, but most of these exceptions are not relevant for price discrimination.

Two exceptions to the processing prohibition could be relevant for online price discrimination. First, a company could legally process special categories of personal data if the data subject gives explicit consent for processing (Art. 9(2)(a) GDPR). But, as noted, it does not seem plausible that many people would give consent to processing for price discrimination.

Second, the prohibition on processing special categories of data does not apply if “the processing relates to personal data which are manifestly made public by the data subject” (Art.

9(2)(e) GDPR). For instance, politicians typically make their political opinions public themselves. Companies could rarely invoke this exception for processing special categories of data for price discrimination. Moreover, even if the prohibition (on processing special categories of data) was lifted, the company would still need a legal basis for the processing. As discussed above, the data subject's consent is generally required for processing personal data for price discrimination. Hence, if a company uses special categories of data for price discrimination and wants to do so in compliance with data protection law, it typically needs the data subject's "explicit consent."

In sum, the data subject's consent seems the only legal basis available in GDPR for personalising prices. In some situations, the ePrivacy Directive *also* requires the customer's consent for personalized pricing, in addition to general data protection law. As noted previously, when companies use cookies and similar techniques, the ePrivacy Directive (Art. 5(3)) also applies, in addition to general data protection law. In such situations, there is another reason why the consumer's consent is required. Cookies can only be legally used when the consumer consents to them (subject to exceptions). If the consumer does not consent (and the company cannot rely on another legal basis), personal data processing is not allowed. Apart from that, the GDPR has specific rules on profiling and on certain types of automated decisions. We turn to those rules in the next section.

Data Protection Law and Automated Decisions

The GDPR contains a specific provision on certain fully automated decisions with far-reaching effects (Art. 22, GDPR). This probably gives another argument why consumer consent is required for personalized pricing. The provision contains an in-principle prohibition of certain fully automated decisions with far-reaching effects and applies, for instance, to automated credit scoring. Article 22 is the successor of Article 15 of the Data Protection Directive, sometimes called the Kafka provision. That Article 15 has not been applied much in practice (Korff 2012).

The main rule of the GDPR's provision on "automated individual decision-making, including profiling" (Art. 22, GDPR) reads as follows:

The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her.

The GDPR defines "profiling" as follows (Art. 4(4)):

'Profiling' means any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, *in particular to analyse or predict aspects* concerning that natural person's performance at work, *economic situation*, health, personal preferences, *interests*, reliability, behaviour, location or movements.¹⁰

The main rule (Art. 22, GDPR) says, in essence, that people may not be subjected to certain automated decisions with far-reaching effects. The GDPR says people have a right "right not to be subject to" certain decisions. But it is often assumed that this right implies an in-principle prohibition of such decisions (De Hert and Gutwirth 2008; Korff 2012; Wachter et al. 2017).

¹⁰ Emphasis added.

Does Article 22 apply to personalized pricing? Slightly rephrasing Mendoza and Bygrave (2017), four conditions must be met for the provision to apply: (i) There is a decision, which is that decision is based (ii) solely (iii) on automated data processing, including profiling; (iv) the decision has legal or similarly significant effects for the person.

With personalized pricing, an algorithm (i) decides, (ii) solely automatically, a price for a customer. Data processed for personalized pricing are (iii) used to evaluate personal aspects of the customer, namely the consumer's willingness to pay and possibly the consumer's economic situation. Therefore, the first three conditions are met when personalized pricing involves purely automated personal data processing.

The fourth condition requires the decision to have "legal effects" or "similarly significantly" affect the person. The Belgian Data Protection Authority (2012) suggests that an advertisement that includes "a reduction and therefore a price offer" has legal effect (Commission for the Protection of Privacy Belgium 2012, par. 80). Presumably, the Authority sees a price offer as an invitation to enter an agreement, which could indeed be seen as having a legal effect. This interpretation would make Article 22 applicable to personalized prices.

An automated decision that similarly significantly affects a person also falls within the scope of Article 22. But, the GDPR does not explain when a decision "significantly" affects a person. Bygrave (2002) (p. 323–324) discusses the predecessor of Article 22 and suggests that personalized pricing, at least when it leads to higher prices, "significantly affects" a person. In the following, we assume that the legal effects criterion is met with price discrimination.

There are exceptions to the in-principle prohibition of certain automated decisions. The prohibition does not apply if the automated decision (i) "is based on the data subject's explicit consent" or (ii) is "necessary for entering into, or performance of, a contract between the data subject and a data controller" (Art. 22(2)).¹¹ As discussed in the previous section, it seems unlikely that many customers would consent to such processing or that companies could invoke necessity for contract performance.

In the unlikely event that a company could rely on the consent or contract exception to bypass the in-principle prohibition, a different rule is triggered: "In the [consent or contract situation], the data controller shall implement suitable measures to safeguard the data subject's rights and freedoms and legitimate interests, at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision" (Art. 22(3)).

In theory, the legal effect of the exception might be as follows. Suppose Fay enters a contract with Nile.com, by clicking "buy this book" on the Nile.com site. Nile.com charges Fay a premium. Nile.com did not ask consent but invokes the contract exception. Later, Fay finds out that she paid a premium. Art. 22 gives her the right to obtain human intervention: the right to speak to a human and to contest the decision. Hence, she could call the helpline and try to convince Nile.com that she should not have paid extra. The situation seems a tad far-fetched, and this right would probably not be of much help to Fay.

If the automated decisions provision applies, the GDPR contains an extra transparency requirement (Art. 13(2)(f) and 14(2)(f)). In addition to the general transparency requirements of the GDPR,

[T]he controller [the company] shall provide the data subject with the following information necessary to ensure fair and transparent processing in respect of the data subject: (...) the existence of automated decision-making, including profiling (...) and, at

¹¹ Another exception (Art. 22(2)(b)) is not relevant for personalized pricing.

least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject.

Hence, in some cases, a company would have to explain that it uses profiling and would have to provide meaningful information about the logic of the profiling process that leads to personalising prices. However, with modern machine learning techniques, it is often difficult to explain the logic behind a decision, when an algorithm, analysing large quantities of data, arrives at that decision (Burrell 2016; Edwards and Veale 2017; Hildebrandt 2015; Kroll et al. 2016).

Some Caveats and Suggestions for Further Research

Above, we argued that European law generally requires transparency and informed consent for personalized pricing. We do not suggest that this is a panacea.

It is too early to conclude that data protection law is the most appropriate instrument to regulate online price discrimination, even when assuming compliance and enforcement to be adequate. For instance, people hardly read, let alone act on information in privacy notices or cookie disclosures (Acquisti and Grossklags 2007). Moreover, some online shops might offer people a take-it-or-leave-it choice. If customers can only buy a product at one shop and that shop informs customers that it personalizes prices, they may have no alternative but to accept the price personalization.

Even if people do not read privacy notices, transparency requirements following from data protection law could be helpful. The transparency requirements could help to unearth how prevalent personalized pricing is and how people respond to transparency about it. Researchers or journalists could read the notices, and name and shame the companies engaging in personalised pricing. Also, regulators could use data protection law's transparency requirements to obtain information about price discrimination. In addition, a data subject could sue a company for breach of data protection rules (Art. 79, GDPR).

Mandated transparency might leave some concerns unresolved, either because it would be insufficiently effective or because some types of price discrimination would be too socially unacceptable to rely on consumer empowerment. If so, regulation that goes further than transparency requirements might be needed.

From a welfare economic perspective, one would ideally only want to allow price discrimination schemes that increase total welfare or, in a more interventionist approach, that increase total consumer surplus. However, such a case-by-case analysis seems an impractical guideline for regulation, because it may lead to lengthy research or litigation, without quick solutions. Some public concerns could be mitigated if the law required that online personalized pricing could only lead to *lower* prices compared to some reference price based on anonymous browsing. This way, personalization could never lead to higher prices compared to this reference price.

But, before considering such interventions, there are still many open questions regarding personalized pricing. More research is welcome on other laws that may limit personalized pricing and which have been left outside the scope of this paper. For instance, if personalized pricing leads (indirectly) to discrimination on the basis of race or gender, it may be prohibited under non-discrimination law (see, e.g., Art. 21, Charter of Fundamental Rights of the European Union; European Agency for Fundamental Rights 2010a, b, p. 29–31). And, under certain circumstances, adapting the price to a customer's nationality may be illegal in the Europe Union (Art 20, 1 Services Directive 2006/123/EC; Schulte-Nölke et al. 2013).

Also, more research is needed into how often online price discrimination occurs and why many people say they dislike it. Would price discrimination be considered fair as long as wealthy people pay more and poorer people less? How can de facto statistical discrimination of protected groups be avoided or dealt with? Are many companies too afraid of consumer backlash to engage in large-scale price discrimination? Or would this be the case if transparency requirements were actually enforced? In 30 years, will we encounter personalized prices all the time?

Conclusion

Online shops are technically capable to offer each website customer a different price, a practice called personalized pricing. An online shop can recognize a customer, for instance through a cookie, and categorize the customer as a price-sensitive or a price-insensitive person. Shops could adapt prices to such profiles, leading to advanced forms of third-degree price discrimination, which could resemble first-degree price discrimination. While such price discrimination could lead to a situation in which large amounts of consumer surplus are extracted to the benefit of companies that apply personalized pricing, the effect on total welfare will be ambiguous. Preliminary evidence on the public attitude towards personalized pricing suggests that people generally dislike it.

This paper argued that European data protection law generally applies to personalized pricing. Data protection law requires a company to inform people about the specific purpose of processing their personal data. And, if a company uses a cookie to recognize somebody, the ePrivacy Directive requires the company to inform the person about the cookie's purpose. This requirement also applies if the purpose is personalising prices. Moreover, data protection law generally requires companies to obtain the consumer's consent for such pricing.

Acknowledgements The authors thank Solon Barocas, Bertin Martens, Akiva Miller, Andrew Odlyzko, Andrew Selbst, participants at the Amsterdam Privacy Conference, EuroCPR and the Privacy Law Scholars Conference, and the anonymous reviewers of the Journal of Consumer Policy for valuable comments on drafts of this paper. All errors remain ours. The paper is partly based on an earlier working paper: Zuiderveen Borgesius, Frederik J., Online Price Discrimination and Data Protection Law (August 28, 2015); Institute for Information Law Research Paper No. 2015-02. <http://ssrn.com/abstract=2652665>.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Acquisti, A., & Grossklags, J. (2007). What can behavioral economics teach us about privacy? In A. Acquisti, S. Gritzalis, C. Lambrinoudakis, & S. di Vimercati (Eds.), *Digital privacy: theory, technologies and practices* (pp.363–377). Boca Raton: Auerbach Publications.
- Amazon (2000). Amazon.com issues statement regarding random price testing. *Amazon News Room*, 27 September 2000. Available at http://phx.corporate-ir.net/phoenix.zhtml?c=176060&p=irol-newsArticle_Print&ID=502821. Accessed 6 June 2017.
- Armstrong, M. (2006). Recent developments in the economics of price discrimination. In R. Blundell, W. Newey, & T. Persson (Eds.), *Advances in Economics and Econometrics, Theory and Applications* (Vol. 2, pp. 97–141). Ninth World Congress. Cambridge: Cambridge University Press.
- Article 29 Working Party (2010). Opinion 2/2010 on Online behavioural advertising (WP 171) 22 June 2010.

- Article 29 Working Party (2012). Appendix of letter to Google (signed by 27 national Data Protection Authorities), 16 October 2012. Available at www.cnll.fr/fileadmin/documents/en/GOOGLE_PRIVACY_POLICY_RECOMMENDATIONS-FINAL-EN.pdf. Accessed 6 June 2017.
- Article 29 Working Party (2013). Opinion 03/2013 on purpose limitation (WP 203), 2 April 2013.
- Article 29 Working Party (2014). Opinion 06/2014 on the notion of legitimate interests of the data controller under Article 7 of Directive 95/46/EC, 9 April 2017.
- Baker, W., Mam, M., & Zawada, C. (2001). Price smarter on the net. *Harvard Business Review*, 79(2), 122–127.
- BBC News (2000). Amazon's old customers 'pay more'. Available at <http://news.bbc.co.uk/2/hi/business/914691.stm>. Accessed 6 June 2017.
- Belgian Data Protection Authority (2012). Opinion of the CPP's accord on the draft regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (Opinion no. 35/2012), unofficial translation (21 November 2012). www.privacycommission.be/sites/privacycommission/files/documents/Opinion_35_2012.pdf. Accessed 26 June 2017.
- Burrell, J. (2016). How the machine 'thinks': understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 1–12.
- Bygrave, L. A. (2002) *Data protection law: approaching its rationale, logic and limits*, (Vol. 10). PhD thesis University of Oslo. Information Law Series, Kluwer Law International.
- CJEU (2008). (ECJ), Case C-524/06 Huber. ECLI:EU:C:2008:724, para. 52.
- CJEU (2016). Case C-582/14 Breyer. ECLI:EU:C:2016:779.
- Commission for the Protection of Privacy Belgium (2012). 'Opinion of the CPP's accord on the draft regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data' (Opinion no. 35/2012), unofficial translation (21 November 2012). www.privacycommission.be/sites/privacycommission/files/documents/Opinion_35_2012.pdf. Accessed 6 June 2017, par. 80.
- De Blasio, B., & Menin, J. (2015). *From cradle to cane: the cost of being a female consumer. A study of gender pricing in new York City*. New York City: Consumer Affairs.
- De Hert P., & Gutwirth S. (2008). Regulating profiling in a democratic constitutional state. In Hildebrandt M., & Gutwirth S. (eds), *Profiling the European Citizen*, Springer.
- ECtHR (1983). Silver and Others v United Kingdom App Nos 5947/72, 6205/73, 7052/75, 7061/75, 7107/75, 7113/75, and 7136/75 (ECHR, 25 March 1983), para. 97.
- Edwards, L., & Veale, M. (2017). Slave to the Algorithm? Why a 'Right to Explanation' is Probably Not the Remedy You are Looking for. <https://ssrn.com/abstract=2972855>. Accessed 6 June 2017.
- European Agency for Fundamental Rights (2010a). *Handbook on European non-discrimination law*. Publications Office of the European Union.
- European Agency for Fundamental Rights (2010b). *Data Protection in the European Union: the Role of National Data Protection Authorities*. Publications Office of the European Union.
- European Commission (2017). Proposal for a Regulation of the European Parliament and of the Council, concerning the respect for private life and the protection of personal data in electronic communications and repealing Directive 2002/58/EC (Regulation on Privacy and Electronic Communications), COM(2017) 10 final. <https://ec.europa.eu/digital-single-market/en/news/proposal-regulation-privacy-and-electronic-communications>. Accessed 6 June 2017.
- Executive Office of the President of the United States (Council of Economic Advisors) (2015). *Big Data and Differential Pricing*. https://www.whitehouse.gov/sites/default/files/whitehouse_files/docs/Big_Data_Report_Nonembargo_v2.pdf. Accessed 6 June 2017.
- Hannak, A., Soeller, G., Lazer, D., Mislove, A., & Wilson, C. (2014). Measuring price discrimination and steering on e-commerce web sites. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, pp. 305–318.
- Hildebrandt, M. (2015). *Smart technologies and the end (s) of law: novel entanglements of law and technology*. Edward Elgar Publishing.
- Korff, D. (2012). Comments on Selected Topics in the Draft EU Data Protection Regulation (17 September 2012). <http://ssrn.com/abstract=2150145>. Accessed 6 June 2017.
- Kroll, J. A., Huey, J., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., et al. (2016). Accountable algorithms. *University of Pennsylvania Law Review*, 165, 633–705.
- Krugman, P. (2000). Reckonings; What Price Fairness? *New York Times*, 4 October 2000. www.nytimes.com/2000/10/04/opinion/reckonings-what-price-fairness.html. Accessed 6 June 2017.
- Kuner, C. (2007). *European data protection law. Corporate Compliance and Regulation*. Oxford University Press.
- Loomes, G., & Sugden, R. (1982). Regret theory: an alternative theory of rational choice under uncertainty. *The Economic Journal*, 92-368, 805–824.
- Mattioli, D. (2012). On Orbitz, Mac Users Steered to Pricier Hotels. *Wall Street Journal*, 23 August 2012.
- Mendoza, I., & Bygrave, L. A. (2017). The Right Not to Be Subject to Automated Decisions Based on Profiling. In T. Synodinou, P. Jougoux, C. Markou, & T. Prastitou (eds.), *EU Internet Law: Regulation and*

- Enforcement. Springer, 2017, Forthcoming; University of Oslo Faculty of Law Research Paper No. 2017–20. <https://ssrn.com/abstract=2964855>. Accessed 6 June 2017.
- Miettinen, T., & Stenbacka, R. (2015). Personalized pricing versus history-based pricing: implications for privacy policy. *Information Economics and Policy*, 33(2015), 56–68.
- Mikians, J., Gyarmati, L., Erramilli, V., & Laoutaris, N. (2013). Crowd-assisted search for price discrimination in e-commerce: first results. In *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies* (pp 1–6).
- Miller, A. A. (2013). What do we worry about when we worry about price discrimination? The law and ethics of using personal information for pricing. *Journal of Technology Law & Policy*, 19, 41–104.
- Narayanan, A. (2013). Online price discrimination: conspicuous by its absence. <http://33bits.org/2013/01/08/online-price-discrimination-conspicuous-by-its-absence>. Accessed 6 June 2017.
- Newell, F. (1997). *The new rules of marketing: how to use one-to-one relationship marketing to be the leader in your industry*. New York: McGraw-Hill.
- Odlyzko, A. (1996). The bumpy road of electronic commerce. In H. Maurer (Ed.), *WebNet 96 – World Conf. Web Soc. Proc.* (pp. 443–456).
- Odlyzko, A. (2003). Privacy, economics, and price discrimination on the Internet. In *Proceedings of the 5th international conference on Electronic commerce* (pp 355–366).
- Odlyzko, A. (2009). Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets. *Review of Network Economics*, 8(1), 40–60.
- Office of Fair Trading (2010). *Online targeting of advertising and prices. A market study*. http://webarchive.nationalarchives.gov.uk/20140402142426/http://www.offt.gov.uk/shared_offt/business_leaflets/659703/OFT1231.pdf. Accessed 6 June 2017.
- Pigou, A. C. (1932). *The economics of welfare*. London: Macmillan & Co..
- Preston McAfee, R. (2008). Price Discrimination, Chapter 20. In *Issues in Competition Law and Policy* (Vol. 1, pp. 465–484). Washington: ABA Book Publishing.
- Schulte-Nölke, H., Zoll, F., Macierzyńska-Franaszczyk, E., Stefan, S., Charlton, S., Barmscheid, M., & Kubela, M. (2013). *Discrimination of Consumers in the Digital Single Market. Study For The Directorate General for Internal Policies Policy Department A: Economic And Scientific Policy (report European Parliament. IP/ A/IMCO/ST/2013–03, PE 507.456'*. [www.europarl.europa.eu/RegData/etudes/etudes/join/2013/507456/IPOL-IMCO_ET\(2013\)507456_EN.pdf](http://www.europarl.europa.eu/RegData/etudes/etudes/join/2013/507456/IPOL-IMCO_ET(2013)507456_EN.pdf). Accessed 6 June 2017.
- Steppe, R. (2017). Prijsdiscriminatie in het digitale tijdperk: Beschouwingen over de nieuwe algemene verordening gegevensbescherming. In M. E. Storme & F. Helsen (Eds.), *Innovatie en disruptie in het economisch recht*. Antwerpen: Intersentia.
- Stigler, G. J. (1987). *Theory of price* (Fourth ed.). New York: Macmillan.
- Strandburg, K. J. (2013). Free fall: the online Market's consumer preference disconnect. *University of Chicago Legal Forum*, 2013(5), 95–172.
- Turow J. (2011). *The Daily You: How the new advertising industry is defining your identity and your worth*. Yale University Press.
- Turow, J., Feldman, L., & Meltzer, K. (2005). *Open to Exploitation: America's Shoppers Online and Offline*. Annenberg Public Policy Center of the University of Pennsylvania. http://repository.upenn.edu/cgi/viewcontent.cgi?article=1035&context=asc_papers. Accessed 6 June 2017.
- Turow J., King, J., Hoofnagle, C. J., Bleakley, A., & Hennessy, M. (2009). Americans Reject Tailored Advertising and Three Activities that Enable it. <http://ssrn.com/abstract=1478214>. Accessed 6 June 2017.
- Valentino-Devries, J., Singer-Vine, J., & Soltani, A. (2012). Websites Vary Prices, Deals Based on Users' Information. *Wall Street Journal*, 23 December 2012. <http://online.wsj.com/article/SB10001424127887323777204578189391813881534.html>. Accessed 6 June 2017.
- Varian, H. R. (1989). Price discrimination. In R. Schmalensee & R. D. Willig (Eds.), *Handbook of industrial organization* (Vol. I, pp. 597–654). Amsterdam: Elsevier.
- Visser, T., Nikiforakis, N., Bielova, N., & Joosen, W. (2014). Crying wolf? on the price discrimination of online airline tickets. In *7th Workshop on Hot Topics in Privacy Enhancing Technologies (HotPETs 2014)*.
- Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *International Data Privacy Law*.
- Zarsky, T. (2002). Mine your own business: making the case for the implications of the data Mining of Personal Information in the forum of public opinion. *Yale Journal of Law and Technology*, 5, 1–56.
- Zarsky, T. (2004). Desperately seeking solutions: using implementation-based solutions for the troubles of information privacy in the age of data mining and the internet society. *Maine Law Review*, 56(1), 13–59.
- Zuiderveen Borgesius, F. J. (2015) *Improving privacy Protection in the Area of Behavioural Targeting*. Kluwer law International.
- Zuiderveen Borgesius, F. J. (2016). Singling out people without knowing their names—behavioural targeting, pseudonymous data, and the new data protection regulation. *Computer Law & Security Review*, 32, 256–271.