



Does Price Personalization Ethically Outperform Unitary Pricing? A Thought Experiment and a Simulation Study

Deni Mazrekaj^{1,2,6} · Mark D. Verhagen^{2,5} · Ajay Kumar³ · Daniel Muzio⁴

Received: 9 March 2023 / Accepted: 28 September 2024 / Published online: 4 November 2024
© The Author(s) 2024

Abstract

Merchants often use personalized pricing: they charge different consumers different prices for the same product. We assess the ethicality of personalized pricing by generalizing and extending an earlier model by Coker and Izaret (Journal of Business Ethics 173:387–398, 2021) who found that price personalization ethically outperforms unitary pricing. Using a simulation analysis, we show that these results crucially depend on the choice of parameters and do not hold universally. We further incorporate additional sources of marginal cost into the utility function that will likely arise from personalized pricing. These include the expectation that personalized pricing is widely considered unfair by consumers who prefer that all consumers are charged the same price (unitary pricing), and that firms often approximate the consumers' willingness-to-pay in ways that may raise negative sentiments among consumers who feel that their privacy is breached. By extending our model with disutility from unfairness perception and disutility from surveillance aversion, we demonstrate that personalized pricing is quickly outperformed by unitary pricing under social welfare functions that tend to prioritize total utility (utilitarianism and prioritarianism), whereas personalized pricing can ethically outperform unitary pricing under social welfare functions that tend to prioritize equality (egalitarianism and leximin). Our findings illustrate various intricacies and dynamics regarding the circumstances under which personalized pricing can be considered ethical.

Keywords Price discrimination · Personalized pricing · Willingness-to-pay · Utility · Artificial intelligence

✉ Deni Mazrekaj
d.mazrekaj@uu.nl

Mark D. Verhagen
mark.verhagen@nuffield.ox.ac.uk

Ajay Kumar
akumar@em-lyon.com

Daniel Muzio
daniel.muzio@york.ac.uk

¹ Department of Sociology, Utrecht University, Utrecht, The Netherlands

² Nuffield College, University of Oxford, New Road, Oxford, UK

³ EMLYON Business School, Lyon, France

⁴ University of York Business School, York, UK

⁵ Leverhulme Centre for Demographic Sciences, University of Oxford, Oxford, UK

⁶ Leuven Economics of Education Research (LEER), KU Leuven, Leuven, Belgium

Introduction

Technological advances have enabled firms to set personalized prices according to customers' characteristics. By applying algorithms on a large collection of customer data, firms can charge different prices to different customers for the same product based on the customers' willingness-to-pay (WTP): the maximum price at which a customer will buy a product (Varian, 1992).¹ The price of headphones Google recommends may be four times higher if you are an affluent buyer than if you are a less affluent buyer according to your search history (Mikians et al., 2012). Similarly, your Uber ride may cost more if your phone battery is low (Martin, 2019), and your hotel prices may be higher if you are booking a room on a Mac than on a Windows laptop (Mattoli, 2012). Personalized pricing has been used throughout different industries and is considered a key tool for increasing firms' profits. This is especially true in industries with a very low marginal cost of production, such as in music and video streaming industries, in which it is very hard to be profitable when prices are set near marginal costs (Coker & Izaret, 2021). For instance, Shiller (2016) estimates that Netflix could increase their profits by almost 15% if they tailored their prices according to customers' web browsing histories.

The ethical understanding of personalized pricing has been a topic of considerable debate. In a systematic review of the literature, Seele et al. (2020) summarize the advantages and disadvantages of personalized pricing. At the consumer level, personalized pricing can lead to an affordable price of products that would have not been available otherwise, price surprises triggering positive emotions, and enhanced personalized services. By contrast, personalized pricing can also lead to privacy loss, potentially higher prices for some products, decreased price certainty, lower price transparency, and an increase in price complexity. The difficulty in assessing the ethicality of personalized pricing can be illustrated by the more ethically ambivalent cases of crucial but unprofitable services, and surge pricing. On the one hand, consider a doctor's office in a rural community in a developing country (Elegido, 2011). This office can only operate profitably if it can charge some customers more than others. In this case, personalized pricing allows treatment of patients that would otherwise not have been able to afford treatment, essentially allowing for progressive redistribution (Bar-Gill, 2019). The alternative to personalized pricing would be no product at all. On the other hand, consider Uber's response to the London Bridge terror attack on 3 June 2017. Following the attack, Uber charged surge

fares as the demand for Uber rides rose substantially (Shead, 2017). Once the company realized what had happened, Uber refunded all the passengers from around the affected area. Thus, although price discrimination may be necessary in some situations to make the provision of a product possible, it may be considered unethical in others.

Previous academic work is largely in disagreement on the ethics of price discrimination. From a perspective that aligns with deontological ethics—under which a valuation based on the current situation is rejected and actions should be morally legitimate regardless of the product of this action—Steinberg (2020) argues that big-data personalized pricing, when employed solely for-profit maximization, fundamentally disrupts the equitable distribution of welfare and resources. He contends that this practice erodes relational equality, not only among consumers but also between consumers and firms, by exacerbating power imbalances in market transactions. His argument is that personalized pricing is “inherently unfair towards consumers, undermining their role as participants in the market by making it unreasonably costly to practice their ability to negotiate prices or shop around.”² This perspective emphasizes the importance of fairness and the protection of individual rights, suggesting that personalized pricing could be morally indefensible if it compromises the fundamental principles of equal treatment and market participation.

A contrasting view is offered by Marcoux (2006) and Elegido (2011) who built their arguments from a more consequentialist rather than deontological position. Marcoux (2006) argued that price discrimination is not more unethical than unitary pricing in which all consumers are charged the same price. His argument is that

“If buyers are subject to a unitary price and some buyers have higher reservation prices than others, it follows that those with higher reservation prices derive greater utility from their purchases than those with lower reservation prices. [...] Consequently, a unitary price affords *unequal degrees of utility enhancement to buyers*—some derive more utility, and others less, when paying the same price.” (Marcoux, 2006, p. 61).

He concluded that “If fairness demands that each buyer enjoys the same welfare from purchasing the same product, then some form of price discrimination [...] is *necessary* to achieve fairness. For *only* a regime of price discrimination could equalize consumer welfare.” (Marcoux, 2006, p. 61). Note that Marcoux did not argue that price discrimination is more ethical than unitary pricing, but rather that it should be preferred if equalizing welfare is what is overall desired.

¹ Although WTP is often also considered to be a range, we follow the traditional approach and consider WTP synonymous with reservation price throughout the paper.

² Note, however, that Steinberg (2020) does not consider Personal Internet Shopping Agents (PISAs) that may effectively eliminate the concern about excessive shopping costs (Lukosius & Hyman, 2018).

Elegido (2011) largely agreed with Marcoux and stated that “price discrimination is a morally neutral practice that businesspeople are entitled to use if it advances their morally legitimate interests.” (Elegido, 2011, p. 639). His argument is that an ethical price is the one which is obtained in an open market that is not characterized by the price of a legal or natural monopoly, even if there are multiple prices paid.

In a recent comprehensive assessment of the ethicality of price personalization, Coker and Izaret (2021) built on the work of Marcoux and Elegido and argued that price personalization is more ethical than unitary pricing from a consequentialist framework. To support their argument, they decomposed the willingness-to-pay into two components: the utility gained from the purchase, and the disutility from money lost from a purchase. Within the context of a stylized example including two types of consumers under certain parameter values (see below), they compared price personalization to unitary pricing using four consequentialist Social Welfare Functions (SWFs)—utilitarian, egalitarian, prioritarian, and leximin—and found that price personalization is preferred to unitary pricing through the lenses of all four SWFs. They concluded that personalized pricing leads to more welfare, even for the consumers.

In this paper, we contribute to the literature by generalizing the model by Coker and Izaret³ and argue that the premise that price personalization is preferred to unitary pricing is not uniformly true. To show this, we extend their original model and expand their results through a simulation analysis using many different parameter values. We show that earlier results stating that personalized pricing outperforms unitary pricing crucially depend on the choice of parameters, like each consumer’s income and the utility gained from the purchase, and do not hold universally. In addition, we argue that the two-component framework should be extended further by including additional sources of marginal cost when comparing unitary versus personalized pricing. Specifically, we decompose the WTP into the two classical components—the utility gained from the purchase and the disutility from money lost for the purchase—and add two additional behavioral components: disutility from possible unfairness perceptions, and disutility from surveillance aversion. We posit that these sources of disutility will likely feature when comparing personalized pricing to unitary pricing but generalize our findings for any source of additional disutility that may arise in personalized pricing. These include sources of disutility that are shared unequally across consumers. We then compare personalized pricing to unitary pricing under the four SWFs considered in earlier work, showing that relatively

small levels of disutility can reverse earlier findings (Coker & Izaret, 2021).

Unfairness Perception and Surveillance Aversion

Previous literature suggests that people’s WTP should be approached through a behavioral lens that moves beyond the classic economic perspective. In this respect, Bar-Gill (2019) took a behavioral approach when he made the willingness-to-pay a function of not only the preferences of consumers, but also of their possible misperceptions. Namely, consumers may have a higher WTP because they misperceive their benefits from the product. For instance, DellaVigna and Malmendier (2006) found that consumers overestimate the benefits from gym subscriptions because they overestimate how often they will go to the gym. In turn, sellers can target this misperception through personalized pricing and charge a price equivalent not to the willingness-to-pay, but to the perceived willingness-to-pay. Bar-Gill (2019) argued that when WTP reflects both preferences and demand-inflating misperceptions, price discrimination hurts consumers and may also reduce market efficiency. We build on this behavioral perspective and argue that the WTP should be decomposed into four components when considering personalized pricing: the utility gained from the purchase, the disutility from money lost for the purchase, the disutility from possible unfairness perceptions, and the disutility from surveillance aversion.

The choice for the two behavioral components is embedded in the theoretical literature on the ethics of algorithms in business (Hermann, 2022; Loi et al., 2022; Seele et al., 2020). As a first behavioral component, we define disutility from unfairness perception as the utility lost because consumers who learn that they paid more than other consumers for the same product may feel wronged. Price fairness is a key component of a pricing system (Xia et al., 2004). Consumers often feel wronged if they are charged a higher price than others as they view it as an unfair or manipulative practice, leading to a decline in demand (Zuiderveen Borgesius & Poort, 2017). Moreover, the use of algorithms to facilitate tailoring prices based on certain characteristics poses ethical challenges regarding fairness (Hermann, 2022). Algorithms may inadvertently perpetuate social disadvantages by segmenting populations. Unfairness can occur when certain customer groups are favored or disadvantaged based on demographic factors. Biased algorithm predictions and unequal treatment can arise from skewed underlying data, including over- or underrepresentation of certain groups, misleading proxy features, or sparse data for specific individuals or groups (Hermann, 2022). These biased predictions can result in unfair and unequal treatment of individuals or groups.

³ This project initially began as a commentary on Coker and Izaret (2021), but it evolved into an original paper, thanks to the valuable feedback and suggestions from the reviewers.

The empirical literature indeed confirms that unfairness perception exists among consumers: many people regard personalized pricing as unfair or manipulative (Anderson & Simester, 2010; Kahneman et al., 1986; Krämer et al., 2018; Li & Jain, 2016; Turow et al., 2005). In a nationally representative survey of 1500 U.S. adults, 76% of the respondents agreed with the statement “It would bother me to learn that other people pay less than I do for the same products.” (Turow et al., 2005). Price discrimination may even be viewed negatively when the customers are advantaged. In the same survey, 72% disagreed with the statement “If a store I shop at frequently charges me lower prices than it charges other people because it wants to keep me as a customer more than it wants to keep them, that’s OK.” The consequences of unfairness perception may be considerable. Using a randomized field experiment of over 50,000 customers, Anderson and Simester (2010) found that customers made fewer subsequent purchases after they found out that the same retailer had been selling the product for less to other customers.

The second utility component we introduce is the disutility from surveillance aversion. Ethical theorists (Culnan & Bies, 2003; Martin, 2016) have recently considered privacy as a social contract: “a mutually beneficial agreement within a community about how information is used and shared” (Martin, 2016, p. 553). Individuals do not give away privacy (as in the access-view of privacy), but they discriminately share information with third parties with certain privacy expectations. Namely, individuals retain the desire to limit the type of information revealed, who has access to information, and how the information is used. Digital surveillance may also be ethically wrong if consumers are coerced to share their data with the seller. For instance, Loi et al. (2022) contend that insurance companies who ask for higher premiums to consumers who want to avoid sharing their data with the insurer is—under certain conditions—an act of coercion by psychological threats. They argue that this form of digital surveillance disrespects a preference for privacy, generates further risks of coercion, and reduces the opportunity for autonomy, authenticity and spontaneity in other choices. Given that personalized pricing relies on an algorithmic estimate of the willingness-to-pay based on customers’ data, a utility loss may occur if personalized pricing suggests to consumers that they are being surveilled by the seller (Miller, 2014; Priester et al., 2020; Rothschild et al., 2019; Turow et al., 2015; Zuiderveen Borgesius & Poort, 2017). Indeed, individuals may not want to share information about themselves, regardless of the consequences, as they value privacy for its own sake (Loi et al., 2022).

Surveying 1012 U.S. consumers between the ages of 18 and 70 who had purchased online in the past six months, Rothschild et al. (2019) found that about 60% of the respondents were concerned about their content consumption,

purchases, online searches, and opt-in wristbands usage being tracked and captured. Similarly, Turow et al. (2015) found that 91% disagreed with the statement “If companies give me a discount, it is a fair exchange for them to collect information about me without my knowing.” When asked to choose about an appropriate punishment for using a person’s information illegally beyond fines, 18% of the surveyed Americans answered that the company should be put out of business and 35% answered that the executives should face jail time (Turow et al., 2009). Moreover, consumers are often unaware of how their data are used by the companies. For instance, Turow et al. (2015) further found that 65% of the respondents did not know that the statement “When a website has a privacy policy, it means the site will not share my information with other websites and companies without my permission” is false. Finally, asymmetric information about parties’ willingness to deal at a specific price point may allow exploitation of lesser-informed parties, providing an additional argument for including disutility from surveillance aversion.

Finally, we want to note that both our proposed sources of disutility may differ depending on the product and market in which personalized pricing is deployed. Disutility may also decrease over time as consumers become more accustomed to personalized pricing. Therefore, although we posit that surveillance and unfairness are likely the most applicable sources of disutility in personalized pricing, we generalize both sources of disutility in our analytical setup to accommodate any source of disutility, as well as different distributions of disutility across consumers. In this way, we, for example, incorporate the possibility that not only consumers who are materially disadvantaged due to personalized pricing may incur unfairness disutility and that consumers who are not directly disadvantaged may also incur negative sentiment from the purchase due to its association with personalized pricing.

Analytical Setup

We illustrate the WTP decomposition starting with the case study of Coker and Izaret (2021) who studied two consumers—Alice (a wealthy businesswoman) and Bob (an undergraduate business student)—who each subscribe to *The Business Journal* and who each have a utility function for income of the form $utility = \ln(income)$. This represents the diminishing marginal utility characteristic, indicating that as income grows, additional utility decreases. Note that there are various ways to define the WTP (Miller et al., 2011), but to maintain comparability to the original findings of Coker and Izaret (2021) we follow their general approach of assuming the utility function of income to follow the natural logarithm of income.

The conventional utility theory defines the indirect utility of no purchase as $u(\text{nopurchase}) = u(\text{income})$. The price paid is zero and there is no utility of consumption. Then, the utility of purchase is $u(\text{purchase}) = u(\text{income} - \text{price}) + u(C)$, with $u(C)$ representing the utility gained from consuming the good. The willingness-to-pay is the price that makes the consumer indifferent between purchasing and not purchasing the good. It follows that

$$u(\text{nopurchase}) = u(\text{purchase})$$

$$u(\text{income}) = u(\text{income} - \text{WTP}) + u(C)$$

$$0 = u(\text{income} - \text{WTP}) - u(\text{income}) + u(C)$$

This implies that the right-hand-side is the consumer surplus or the net utility surplus amounts to zero at WTP. Coker and Izaret (2021) decomposed the willingness-to-pay into two components: utility gained from the purchase— $u(G)$ —and disutility from money lost from a purchase— $u(L)$.

In the case of Coker and Izaret (2021), they proceeded by setting Alice's average income to \$25,000, and Bob's to \$2500 per month; their corresponding utilities from income are then 10.1 utils and 7.8 utils. Further, they assumed that Alice gets 0.1 utils from her subscription because she needs the journal subscription to do her job, whereas Bob gets 0.05 utils as he only reads it as a hobby. Two scenarios are considered. The first scenario corresponds to unitary pricing in which *The Business Journal* charges the same price of \$20 to both Alice and Bob. The second scenario corresponds to personalized pricing in which *The Business Journal* uses an artificial intelligence (AI) algorithm to estimate each consumer's WTP and charges a higher price to the consumer with the larger WTP. Specifically, Alice is charged \$35 for the journal subscription and Bob only \$5.⁴

We start our exposition by generalizing the case studied by Coker and Izaret (2021) outlined above, noting that their choices of parameter estimates can be viewed as arbitrary and may not be representative of real-life settings. For instance, it is unlikely that Alice is charged 700% more than Bob due to price personalization in many industries. It also may or may not be realistic that Alice's utility of the product is higher than that of Bob. Their setup can be generalized as a function of three input parameters per consumer: their income, the unit price of the product, and the personal utility gained from the purchase. In our main results, we provide

simulations illustrating the welfare dynamics when varying the assumptions around income, price, and personal utility from the product for both Bob and Alice.

We then extend the original Coker and Izaret (2021) decomposition with two additional behavioral components: disutility from unfairness perception— $u(F)$ —and disutility from surveillance aversion— $u(S)$. The WTP can be formulated as follows:

$$\text{WTP} = u(G) - u(L) - u(F) - u(S)$$

or

$$\text{WTP} = u(G) - u(\text{LFS})$$

$$\text{with } u(\text{LFS}) = u(L) + u(F) + u(S)$$

A rational individual would buy the product if the utility gained from the purchase $u(G)$ is larger or equal to the sum of the three disutility components $u(\text{LFS})$, and would not buy the product if $u(G) < u(\text{LFS})$.

In the first scenario that corresponds to unitary pricing, the WTP decomposition includes only the two original components, namely the utility gained from the purchase $u(G)$, and disutility from money lost from a purchase $u(L)$. Given that both Alice and Bob are charged the same price, there is no disutility from unfairness perception $u(F)$ nor from surveillance aversion $u(S)$. We now extend this example and assume that both Alice and Bob are aware that *The Business Journal* uses algorithmic personalized pricing and that both Alice and Bob have an aversion to their information being used for the algorithm ($u(S) > 0$), and Alice perceives unfairness by being charged a higher price than Bob ($u(F) > 0$). Contrary to the unitary pricing scenario, the WTP decomposition now includes all four components: utility gained from the purchase $u(G)$, disutility from money lost from the purchase $u(L)$, disutility from unfairness perception $u(F)$, and disutility from surveillance aversion $u(S)$.

To illustrate the calculations behind our simulation analysis under the four social welfare functions, in what follows, we first present a simple thought experiment of Alice and Bob as in Coker and Izaret (2021). For this purpose, we take a conservative approach, and assume that Alice's disutilities $u(F)$ and $u(S)$ are each one twentieth of $u(G)$ and thus a proportion of the utility gained from the product. We base these estimates on recent lab experiments assessing the impact of price discrimination on consumer behavior (Allender et al., 2021).⁵ Across these examples, producer surplus is equal across unitary and personalized pricing, as long as

⁴ Note that as a consumer's WTP is measured algorithmically by the firm, the estimated WTP will naturally be noisy. This is a result of both the imperfect accuracy of the algorithm, but also of the consumers' uncertainty about their own WTP. Consequently, price-setting will likely be in the direction of the true WTP but not equal to each individual's WTP, as is the case in our example (see Discussion section).

⁵ We base our 5% estimate of disutility due to unfairness perception on lab results indicating that a consumer's utility decreases as a function of the differences in prices charged to consumers. In general, it is difficult to transfer findings from the empirical literature to our context, as both the empirical context as well as the controls can vary. However, a relatively small discount seems appropriate.

both transactions occur—i.e., $u(G) \geq u(LFS)$. We then relax these assumptions in the simulations. We further note that our example introduces the possibility for imbalanced disutility due to personalized pricing: surveillance loss is shared by both Alice and Bob, but only Alice receives unfairness disutility. As we reduce unfairness loss, we illustrate the case where disutility is completely symmetrical across both consumers. As we reduce surveillance disutility, we approach a scenario where only one consumer suffers disutility from personalized pricing. This flexibility allows many different types of disutility to be assessed that move beyond surveillance and unfairness.

Social Welfare Functions

To ethically evaluate unitary and personalized pricing under our conservative estimates of disutility stemming from unfairness perception and surveillance aversion, we employ the concept of Social Welfare Functions (SWFs). Introduced by Bergson (1938), and popularized by Arrow (1951), SWFs are functions that rank states according to their social desirableness. SWFs are widely used in welfare economics and assume that the overall desirability of a state can be described by a single number, in our example the total utility surplus of Alice and Bob. We follow Coker and Izaret (2021) and invoke four consequentialist SWFs: utilitarianism, egalitarianism, prioritarianism, and leximin that relate both players' utility from the transaction to an overall welfare indication. Note that many alternative SWFs other than those presented above exist. We focus solely on the four consequentialist SWFs to enable direct comparisons with Coker and Izaret (2021). The debate regarding which ethical framework is favored goes beyond the scope of this paper. In what follows, we assess the impact of our two behavioral components when added to the original example by Coker and Izaret (2021).

Utilitarianism

From a utilitarian perspective (Bentham, 1780), the preferred pricing would be the one in which the sum of the total net utility surpluses is maximized. Thus, we calculate the total net utility surpluses for both Alice and Bob under both unitary and personalized pricing using (1):

$$SWF_{utilitarianism} = usurplus_{Alice} + usurplus_{Bob}, \quad (1)$$

where $usurplus_{Alice(Bob)}$ is the net utility surplus obtained by Alice (Bob) from the transaction.

Unitary pricing:

$$usurplus_{Alice} = 0.1 - [\ln(25,000) - \ln(24,980)] = 0.0992$$

$$usurplus_{Bob} = 0.05 - [\ln(2,500) - \ln(2,480)] = 0.0420$$

$$SWF_{utilitarianism} = 0.0992 + 0.0420 = 0.1412$$

Personalized pricing:

$$\begin{aligned} usurplus_{Alice} &= 0.1 - [\ln(25,000) - \ln(24,965)] \\ &\quad - 0.1 * \frac{2}{20} = 0.0886 \end{aligned}$$

$$\begin{aligned} usurplus_{Bob} &= 0.05 - [\ln(2,500) - \ln(2,495)] \\ &\quad - 0.05 * \frac{1}{20} = 0.0455 \end{aligned}$$

$$SWF_{utilitarianism} = 0.0886 + 0.0455 = 0.1341$$

Given that the total welfare⁶ is higher under unitary pricing than under personalized pricing, unitary pricing is favored under a utilitarian SWF from a consumer's perspective. This is contrary to the welfare found by Coker and Izaret (2021) when excluding these behavioral components. Three points about this calculation deserve special attention. First, we assumed that both Alice and Bob have an aversion to their information being used for the algorithm, as we believe surveillance aversion is unrelated to whether a person is favored by the pricing scheme or not. However, removing this restriction hardly changes the total net utility surplus under personalized pricing, as discussed in more detail below. Second, we assumed that the disutility of unfairness perception and surveillance aversion is one twentieth of the utility gained from the purchase. Although we believe this is a rather conservative estimate, we find that setting each disutility component (unfairness perception and surveillance aversion) to 2.16% of the utility gained from the purchase would already make utility pricing equivalent to personalized pricing under the utilitarian SWF in this specific context. Thus, even very conservative estimates of the disutilities are sufficient to render unitary pricing the favored pricing scheme.

It should be noted that this analysis does not apply to industries that would not exist under unitary pricing. As highlighted in the introduction, the example of a doctor's office in a rural community in a developing country (Elegido, 2011) helps illustrate that some industries can only operate profitably if they charge some customers more than others. Under the utilitarian SWF, personalized pricing would be favored in this case as it produces a surplus, whereas unitary pricing does not. Said otherwise, personalized pricing would

⁶ Although, strictly speaking, welfare is not the same as surplus, we use the word welfare throughout the paper for ease of interpretation in line with our use of the social welfare functions.

be Pareto-improving as everyone would be better off and no one would be worse off. This situation is beyond the scope of our example, as both transactions occur across unitary as well as personalized pricing.

In sum, our calculation shows that unitary pricing can already ethically outperform personalized pricing in most industries that do not depend on personalized pricing for existence for small levels of disutility due to unfairness perception or surveillance aversion. This is not surprising, as both sources of loss directly decrease both players' net utility surplus and utilitarian welfare is a simple sum of both. Combined with the fact that the gain in welfare due to personalized pricing was already fairly small to begin with leads to a reversal of the earlier results of Coker and Izaret (2021).

Egalitarianism

From an egalitarian perspective, it is not the total utility surplus that is of primary interest, but the distribution of the individual utility surpluses. Whereas utilitarianism does not take equality into consideration, egalitarianism prioritizes equality above all else and does not consider the total utility surplus. Hence, egalitarianism favors the pricing scheme under which the utility surpluses are closer to equality:

$$SWF_{\text{egalitarianism}} = |(usurplus_{\text{Alice}} - usurplus_{\text{Bob}})| \quad (2)$$

In our example, the utility surplus gap between Alice and Bob is 0.0572 (0.0992–0.0420) under unitary pricing, and 0.0431 (0.0886–0.0455) under personalized pricing. The utility surplus gap is narrowed from 0.0572 to 0.0431 under personalized pricing, an improvement of 24.66%. Unlike under utilitarian SWF, personalized pricing is favored over unitary pricing under the egalitarian SWF. Effectively, the penalty on Alice's utility surplus due to a higher unit price, as well as the combination of unfairness perception and surveillance aversion means that the difference between Alice's and Bob's utility surplus from the product narrows under personalized pricing. As discussed in more detail below, this is primarily because of the considerable differences in the baseline utility gained from the product as well as income differences between Alice and Bob. This finding echoes the viewpoint from Marcoux (2006) that personalized pricing can address differences in reservation prices.

Prioritarianism

Whereas utilitarianism is criticized based on a lack of equality considerations, egalitarianism is criticized based on the 'Levelling Down Objection' (Temkin, 2000). This objection states that egalitarianism can reject Pareto-improving states as it does not consider the total utility surplus. For

instance, egalitarianism would prefer Alice and Bob both having a utility surplus of 0.1 utils to Alice having 0.1 utils and Bob 0.2 utils, even though Bob is now better off, while Alice is not worse off. Prioritarianism (Parfit, 1998) aims to reconcile both utilitarianism and egalitarianism by considering both the total utility surplus and equality. Specifically, prioritarianism gives extra weight to worse-off individuals, by weighing their utility based on how much utility an individual already has. For instance, if Alice moves from 0.8 to 0.9 utils, whereas Bob moves from 0.1 to 0.2 utils, Bob will receive a larger weight than Alice. In essence, prioritarianism calculates the priority-weighted change in utilities before and after buying the product for each individual:

$$SWF_{\text{prioritarianism}} = \ln\left(\frac{\ln(\text{income})_{\text{Alice}} - usurplus_{\text{Alice}}}{\ln(\text{income})_{\text{Alice}}}\right) + \ln\left(\frac{\ln(\text{income})_{\text{Bob}} - usurplus_{\text{Bob}}}{\ln(\text{income})_{\text{Bob}}}\right) \quad (3)$$

In our example, Alice had a utility of $\ln(25,000) = 10.1266$ and Bob $\ln(2500) = 7.8240$ before buying *The Business Journal* subscription. After buying the subscription, Alice gained 0.0992 under unitary pricing and 0.0886 under personalized pricing, whereas Bob gained 0.0420 and 0.0455 under unitary and personalized pricing, respectively (see the utilitarian SWF). The priority-weighted change in utility can then be calculated using (3):

Unitary pricing:

$$\begin{aligned} \text{priority-weighted utility surplus}_{\text{Alice}} &= \ln\left(\frac{10.1266 + 0.0992}{10.1266}\right) = 0.0097 \\ \text{priority-weighted utility surplus}_{\text{Bob}} &= \ln\left(\frac{7.8240 + 0.0420}{7.8240}\right) = 0.0054 \end{aligned}$$

$$\begin{aligned} \text{priority-weighted utility surplus}_{\text{Bob}} &= \ln\left(\frac{7.8240 + 0.0420}{7.8240}\right) = 0.0054 \end{aligned}$$

$$SWF_{\text{prioritarianism}} = 0.0097 + 0.0054 = 0.0151$$

Personalized pricing:

$$\begin{aligned} \text{priority-weighted utility surplus}_{\text{Alice}} &= \ln\left(\frac{10.1266 + 0.0886}{10.1266}\right) = 0.0087 \\ \text{priority-weighted utility surplus}_{\text{Bob}} &= \ln\left(\frac{7.8240 + 0.0455}{7.8240}\right) = 0.0058 \end{aligned}$$

$$\begin{aligned} \text{priority-weighted utility surplus}_{\text{Bob}} &= \ln\left(\frac{7.8240 + 0.0455}{7.8240}\right) = 0.0058 \end{aligned}$$

$$SWF_{\text{prioritarianism}} = 0.0087 + 0.0058 = 0.0145$$

It appears that, under unitary pricing, the total priority-weighted utility surplus is higher than under personalized

pricing. Thus, unitary pricing is favored under a prioritarian SWF for relatively low levels of disutility due to unfairness perception and surveillance aversion. So, although the utility surpluses of both Alice and Bob lie closer together (as we saw under egalitarianism), the overall loss in utility surplus still renders personalized pricing inferior to unitary pricing under a prioritarianist SWF.

Leximin

The leximin SWF (Sen, 1979) states, in our context, that the pricing that maximizes the utility surplus of the least well-off individual is the most desirable pricing. If the least well-off individual is equally well off under the two pricing regimes, then the next least well-off takes lexical priority:

$$SWF_{\text{leximin}} = \min(usurplus_{\text{Alice}}, usurplus_{\text{Bob}}) \quad (4)$$

For instance, if Bob (who is least well off) is equally well off under unitary and personalized pricing, we would look at Alice's utility under the two regimes. This distinguishes leximin from maximin (Rawls, 1971), whereby the latter is more ambiguous, as it does not include this tie-breaking condition.

Under unitary pricing, the least well-off individual, Bob, obtains a utility surplus of 0.0420 from the purchase, whereas under personalized pricing his utility surplus increased to 0.0455. Given that Bob's utility surplus increased, and Bob is least well off under the transaction, personalized pricing is favored under the leximin or maximin SWFs. Nonetheless, it is useful to calculate how large Bob's surveillance aversion would have to be to render unitary pricing as the preferred option. Note that, under the leximin criterion, unitary pricing would be preferred over personalized pricing if Bob's utility surpluses under the two pricing regimes would be equal. This is because we would look to the next least well-off individual—Alice, who has a higher utility surplus under unitary pricing than under personalized pricing. We calculate that surveillance aversion would have to be about one-eighth of the utility gained from the purchase to render unitary pricing the ethically preferred pricing scheme under the leximin criterion. Thus, it does not take an extreme surveillance aversion to reverse the results. More generally, leximin prefers personalized pricing if the least well-off individual's improvement in utility surplus due to a lower price is offset by any loss due to surveillance.

Simulation Results

Having illustrated our calculations for a specific choice of parameters, we now present our main analysis by expanding the above exposition in two ways. First, we critically assess the parameter assumptions by Coker and Izaret (2021) that

informed their original claim that personalized pricing is preferred over unitary pricing by using our more general framework in which we vary the starting values for income, personal utility, and unit prices. Then, we further assess the implications of varying our own additional components of unfairness perception and surveillance aversion. In what follows, we start with a brief discussion of utility dynamics, after which we assess the implications of personalized pricing using the SWFs.

Utility Dynamics

The various SWFs considered above depend crucially on the utility surpluses of Alice and Bob: $usurplus_{\text{Alice}}$ and $usurplus_{\text{Bob}}$. Each individual's utility surplus is composed of the utility gained from the product minus the disutility from the purchase. The latter solely depends on the disposable income available to the individual and the unit price charged to them. These dynamics are illustrated for various unit prices and income levels in Fig. 1 in the left panel. The assumptions by Coker and Izaret (2021) are illustrated in this graph by markers A and B. On the other hand, the utility surplus from the product solely depends on the assumed personal utility gain from the product and any loss that reduces this utility, in our case loss related to unfairness perception and surveillance aversion. The dynamics for both Alice and Bob reflect how this utility surplus declines as a function of such loss under personalized pricing in Fig. 1 in the right panel. Note that it is assumed here that Alice is charged a higher price than Bob and thus incurs disutility from both unfairness perception and surveillance aversion. The Coker and Izaret (2021) assumptions (with unfairness perception and surveillance aversion set at 0%) are again illustrated by markers A and B.

Some important dynamics emerge from Fig. 1. First, as income rises, the relative disutility from being charged a higher unit price becomes considerably less severe as can be observed by comparing the curves for Alice and Bob, on the left. Second, assumptions regarding the initial utility from the product play a large role in determining each person's overall utility from the purchase. Third, disutility from unfairness perception, which is only relevant in our setup for one of the two persons under personalized pricing, can introduce an inflection point where Alice's utility from the product in fact becomes lower than that from Bob. These dynamics have several implications that we discuss further during our simulations.

Simulation setup

In what follows, we present results when varying four of the eight possible parameters that feature in the calculation of each SWF: both players' utility from the product, both players' income, the two prices charged to both players under




Fig. 1 Disutility from acquisition of the product as a function of one's disposable income and the unit price charged (left). Utility from acquisition of the product as a function of one's personal utility from the product (y-axis) and loss due to unfairness perception and / or surveillance aversion. Note that it is assumed that player A is charged a higher price and thus incurs disutility from both unfairness percep-

tion and surveillance aversion, whereas player B only incurs disutility from surveillance aversion. The assumptions on player A and B by Coker and Izaret (2021) are visualized in both graphs by markers A and B—note that neither player incurs disutility from unfairness perception or surveillance aversion in their setup

personalized pricing, and, finally, the disutility incurred to both players due to personalized pricing. Note that the number of combinations of the parameters scales exponentially and we are precluded from assessing all the possible combinations. Therefore, we have made an online tool where the reader can evaluate the parameter input of their choice at https://bright-box.shinyapps.io/price_discrimination/. We also provide the code to our simulations at: https://github.com/MarkDVerhagen/price_discrimination_tool/.

Across each simulation, we present results for three combinations of incomes [$income_{Alice}, income_{Bob}$]: ([25000, 2500], [13750, 13750], [2500, 25000]). The first concerns the original setup by Coker and Izaret (2021), where the individual with the higher personal utility from the product also has the higher income. The latter two illustrate cases where incomes are equal and where the higher personal utility player has the lower income, respectively. Besides these three income settings, we further vary two out of the six remaining parameters per simulation and show the difference in welfare between personalized and unitary pricing. Note further that for simplicity, we constrain our simulations from the perspective of the producer. Specifically, when varying the prices charged to Alice and Bob, the sum is always equal to the sum charged under unitary prices. Furthermore, we only consider parameter settings where both transactions occur—i.e., $u(G) \geq u(LFS)$). In other words, producer surplus is equal throughout the simulations. Note further that

the equal sum constraint on prices reduces the two price parameters to one, as they are directly related.

In our first simulation, we disregard disutility from unfairness perception and surveillance aversion but relax the assumptions by Coker and Izaret (2021) on starting income, personal utility from the purchase, and unit prices charged in terms of the four welfare functions considered. We consider this setup to be the generalization of their original case. Figure 2 illustrates how personalized pricing performs relative to unitary pricing when varying the personal utility of both Alice and Bob. Note that we cap the minimum utility for both, such that the transaction would never yield negative utility. As can be seen from the middle column, which maintains the income assumptions by Coker and Izaret (2021), both utilitarianism as well as prioritarianism improve welfare regardless of starting values of the personal utility of both players. This is because the richer individual (player A) gets charged more than the poorer one (player B) combined with the income-price dynamics from Fig. 1. Conversely, setting starting income equal or reversing them would lead to consistent negative welfare due to personalized pricing for utilitarianism; for prioritarianism this holds under reversed income. Thus, if pricing were to be personalized based on personal utility rather than personal income there can be welfare loss due to personalization. For egalitarianism and leximin, negative welfare requires that the person being charged more has lower utility from the product than




Fig. 2 Welfare under the Coker and Izaret (2021) setup without disutility from unfairness perception and surveillance aversion when varying the personal utility for both persons. Each column further

represents different starting values for the disposable income of both players (Coker and Izaret (2021) income assumption in the middle column). Each row reflects a different welfare function





Fig. 3 Welfare under the Coker and Izaret (2021) setup without disutility from unfairness perception and surveillance aversion when varying the price charged to both persons. Each column further rep-

resents different starting values for the disposable income of both players (Coker and Izaret (2021) income assumption in the middle column). Each row reflects a different welfare function

Fig. 4 Difference in welfare under personalized pricing versus unitary pricing for various levels of unfairness perception and surveillance aversion under the Coker and Izaret (2021) setup. Note that the player being charged more incurs disutility from unfairness perception in addition to disutility from surveillance aversion. The latter is incurred by both players



the person being charged less under the Coker and Izaret (2021) assumptions for income. This could happen if pricing is done solely on disposable income and personal utility is higher for the poorer individual. In the Coker and Izaret (2021) assumptions, higher income and higher personal utility align, ignoring situations where income inequality is not consistent with utility inequality. In other words, the alignment of higher utility and income is a crucial condition for their main results.

In terms of varying the unit price charged to both players, similar dynamics arise as presented in Fig. 3. Regardless of what exact prices are being charged, under the assumptions of Coker and Izaret (2021) that the wealthier individual has the higher personal utility, every combination of prices that sum to 40 and where the person with higher utility is charged more leads to welfare gains for every welfare function. This is no longer the case when switching income levels for utilitarianism and prioritarianism (and welfare is equal for these two cases under equal incomes).

Thus, if Alice had the higher utility from the product but lower income than Bob, personalized pricing would be detrimental to welfare. These findings place important conditions on earlier findings like those of Coker and Izaret (2021) who ignored asymmetries in utility from the product and starting income. However, there are many cases where utility for the same product would intuitively be higher for low-income individuals relative to their high income counterparts. Note furthermore that these dynamics also illustrate how even under the Coker and Izaret (2021) assumptions of income and personal utility, there is scope to increase producer welfare while still improving consumer welfare relative to unitary pricing—i.e., the price point can be moved away

from the diagonal line where prices sum to 40, for example charging a price of 36 to Alice and 4.5 to Bob.⁷

We now turn to the discussion of disutility from unfairness perception and surveillance aversion in Fig. 4. The player being charged more incurs disutility from unfairness perception in addition to disutility from surveillance aversion. The latter is incurred by both players. As we vary the loss incurred beyond the 5% assumed above, we observe an inflection point for egalitarianism and leximin. This is intuitive as egalitarianism considers the difference in utility gained and these utilities converge toward one another as loss increases until, at the inflection point, the utility of Alice starts to dip under that of Bob, thus decreasing egalitarian welfare again. Similarly, leximin considers the utility gained by the worst off. This is Bob until the inflection point, after which the slope doubles because Alice incurs both sources of loss. Note that the Coker and Izaret (2021) welfare calculations can be found when considering 0% on the x-axis in Fig. 4. We further illustrate the 5% disutility finding which is based on empirical estimates of disutility due to unfairness perception with the vertical line.

In our earlier exposition, we specified a particular parametrization in which unfairness perception and surveillance aversion affect the utility of Alice and Bob as a proportion of the utility of the product. We next relax this parametrization of both sources of loss. Specifically, we simply define an overall loss pertaining to Alice, l_{Alice} , and Bob, l_{Bob} , without

⁷ We provide further dynamics when varying both the personal utility and unit price charged to one player while retaining the Coker and Izaret (2021) assumptions for the other in Figs. 6 and 7 in the Appendix.

further assumptions. In other words, we study the full range of possibilities in which Alice or Bob is affected by both sources of loss. This covers concerns that either of the two types of loss follow a non-linear or step-wise design as one might expect based on for instance Prospect Theory that suggests that individuals may exhibit risk aversion for positive prospects and risk seeking for negative ones (Kahneman & Tversky, 1979)⁸ or cases where both Alice and Bob experience disutility due to unfairness perception.⁹ Making both player's loss non-parametric also covers instances where unfairness perception is a function of the differences between the prices charged to both (Allender et al., 2021), or might vary by characteristics of the buyer (Bolton et al., 2010) and covers cases where there is a *higher* loss for Bob even though they should not formally incur a loss due to unfairness perception. This could be, for example, if surveillance loss is a function of disposable income or some absolute level of utility. Finally, a non-parametric setup also covers the possibility that there might be heterogeneity because of cultural, time, or market-related differences in fairness or surveillance loss, for instance when norms differ among actors as might be the case in settings where bidding or haggling is more common to economic transactions. Given these varied determinants of disutility, it seems reasonable to simply consider all possible levels of loss up to the point that the transaction no longer occurs—i.e., $u(G) \geq u(LFS)$.¹⁰ It also generalizes our results to any type of disutility that may be incurred beyond surveillance and unfairness disutility per se. For tractability, we otherwise maintain the Coker and Izaret (2021) assumptions noting that the dynamics we mentioned in Figs. 2 and 3 will affect the inflection point at which personalized pricing starts underperforming unitary pricing.

The results comparing personalized pricing versus unitary pricing as a function of arbitrary levels of loss are found in Fig. 5. Both the full dynamics as well as those following

our earlier specification (dashed line) indicate that increased loss due to unfairness perception and surveillance aversion can negate any positive welfare gains quite rapidly. In other words, minor disutility is often sufficient to nullify any gain in welfare. This holds for most welfare functions except egalitarianism which instead improves strongest whenever the loss of Alice is higher than that of Bob until the inflection point in utility is reached. In these graphs, we highlight the assumed levels of disutility of 5% which we based on lab results assessing disutility from unfairness (Allender et al., 2021), and only include loss for which both transactions still occur.

Several findings arise from these results. First, as already became clear earlier, comparatively small levels of utility loss can negate any welfare gains due to personalized pricing. Often, this makes intuitive sense as there is a general loss of utility affecting the consumers. In particular, for SWFs that emphasize overall utility (utilitarianism and prioritarianism). Second, for egalitarianism and leximin, there is a considerably larger domain in which welfare gains can be maintained under personalized pricing. This is particularly the case when loss is centered around the player with the higher income, as well as the higher utility from the product. From an egalitarian perspective, welfare gains are strongest around the point where loss is such that any differences in income and utility from the product are negated. In other words, when personalized pricing leads to disutility from unfairness perception and surveillance aversion, it can be an effective tool to equate utility across consumers, although only by virtue of reducing overall utility and doing so unequally. Third, we include likely estimates based on the empirical literature, which suggested limited loss when expressed as a percentage of the utility gained from a product. Our results show that in reality, personalized pricing could very well be beneficial across some SWFs but detrimental in others and will likely depend on small nuances concerning the context within which it is implemented. However, there is clearly limited scope for general statements as to the dominance of personalized pricing over unitary pricing when taking our two behavioral components into account.

Limitations and Assumptions

The rationale for our paper relies on three assumptions underlying the results presented above. These are (i) that both disutility due to unfairness perception as well as due to surveillance aversion are reasonable, (ii) that consumers know whenever firms are implementing price personalization, and (iii) that personalized pricing does not lead to complete equation of price-setting to WTP. Before providing a general discussion of our findings, we briefly touch on these assumptions below.

⁸ This so-called reflection effect may lead individuals to weigh perceived losses more than perceived gains. In our setting, prospect theory would imply that consumers may weigh the disutility disproportionately to the positive utility. A similar argument concerns how personal utility may be perceived by an individual. Utility can be personal or internalized impersonal. Think of the latter as feeling good about benefitting society, such as supporting sustainability. Disutility, on the other hand, is inherently personal. This may add to weighting disutility disproportionately to the positive utility.

⁹ As mentioned in the introduction, Turow, Feldman, and Meltzer (2005) found that 72% of the U.S. respondents disagreed with the statement “If a store I shop at frequently charges me lower prices than it charges other people because it wants to keep me as a customer more than it wants to keep them, that’s OK.” Thus, it is possible that Bob also experiences disutility from unfairness perception, possibly in a milder manner than Alice.

¹⁰ As the financial cost of the product is small relative to both Alice’s and Bob’s income, this effectively means assessing loss up to a complete negation of utility, i.e., 100% loss due to the combination of unfairness perception and surveillance aversion.




Fig. 5 Difference in welfare under personalized pricing versus unitary pricing for any level of loss incurred by player A and player B expressed as a percentage of the utility from the purchase. Dashed line reflects our initial specification of additive proportional loss

Unfairness Perception and Surveillance Aversion

Most previous studies have not taken unfairness perception into account, although the empirical literature persistently found that people generally regard personalized pricing as unfair or manipulative as it violates the ‘Equal Treatment Norm’ of market exchanges (Anderson & Simester, 2010; Kahneman et al., 1986; Krämer et al., 2018; Li & Jain, 2016; Turow et al., 2005). The argument behind the exclusion of unfairness perception is that the equal treatment norm may either not apply to pricing issues (Elegido, 2011), or that even if it does, an equal treatment norm would more plausibly apply to utility than prices themselves as the “currency of egalitarian justice is utility” (Coker & Izaret, 2021, p. 8). However, whether the consumers’ reaction to personalized pricing is philosophically plausible or not, the empirical literature does indicate that disutility from unfairness perception exists (Allender et al., 2021). Thus, even if the reaction may appear irrational, it should still be included as a disutility component when assessing the desirability of different pricing schemes.

Another objection to including unfairness perception is that firms can use marketing strategies to influence unfairness perception. It is exactly this argument that precluded Coker and Izaret (2021) from including unfairness perception in their original two-components WTP decomposition. They argue that firms can convince consumers by highlighting that some people with a higher WTP are paying more than them. However, Turow et al. (2005) show that consumers view price personalization as unfair even if they are advantaged by it. Moreover, it is unlikely that marketing strategies can effectively eliminate unfairness perception for the consumers who feel that they are at the top of the price distribution and are paying substantially more than most people.

Another argument is that price personalization will be perceived as less unfair as it is used more widely and as it becomes more familiar. To some extent, this disutility could be ephemeral and dissipate when norms change. Indeed, things that are considered harmful at one point in time may no longer bother society in the future. If we believe that the disutility of unfairness may change over time, then it can be argued that this disutility should not be incorporated into policy decisions. However, this is not necessarily the case as the drawbacks of price personalization also become more known over time. As shown by Turow, et al. (2015), many people are not yet aware that price discrimination occurs. Therefore, we would argue that unfairness perception should not be excluded from the WTP decomposition.

An objection to targeting surveillance aversion may be that it is not relevant as many countries passed regulations on data protection. For instance, the European Economic Area passed the General Data Protection Regulation (GDPR) in

2016 that was implemented in 2018. Many U.S. states have also implemented a similar framework, the California Consumer Privacy Act being the first in 2018. In essence, these laws aim to protect consumers by ensuring that personal data processing happens fairly, lawfully, and transparently. For instance, the GDPR implies that firms must inform consumers if they personalize prices and consumers must give consent to their personal data being used (Zuiderveen Borgesius & Poort, 2017). If a company uses e-cookies to recognize a recurrent consumer, the ePrivacy Directive requires the firm to inform the person about the cookie’s purpose and to ask for consent. Although helpful, it is unlikely that these types of regulations eradicate surveillance aversion even when compliance is substantial. For instance, Acquisti and Grossklags (2008) found that people hardly read information in privacy notices or cookie disclosures. Even if they do read the information, people tend not to act on this information. Thus, surveillance aversion is likely to exist even in the presence of personal data processing regulations.

Another viewpoint could be that as markets become increasingly digital, most consumers are aware and accept that they are being surveilled regardless. In other words, surveillance loss is a part of market participation rather than a characteristic of a specific product. Such convergence may, in the future, reduce the extent to which surveillance loss is actively experienced by the consumer. This convergence toward surveillance being completely associated with market participation will likely be limited by the implementation of before-mentioned regulation. Specifically, by requiring firms to disclose the collection of personalized data, it becomes possible for firms to differentiate from one another in terms of the *extent to which* they surveil their consumers (e.g., by explicitly stating that personalized data are never used for commercial purposes). Without any regulatory framework, such differences may become less feasible within a competitive marketing setting. Additionally, consumers are still likely to actively experience surveillance loss even though large parts of retail have moved online. However, the possibility that this might change in the future and surveillance loss might become a part of market participation has to be taken into account.

It is possible to argue, especially under the utilitarian SWFs as mentioned previously, that personalized pricing should be preferred over unitary pricing in industries that would not exist if they were not able to charge higher prices to some consumers. For instance, Coker and Izaret (2021) posit that it may be beneficial to use a personalized pricing scheme for spinal muscular atrophy medication, for which a recently approved gene therapy costs \$2.1 million. On the other side of the coin, personalized pricing may lead to price gouging, exploitation, and false advertising in some industries (Snyder, 2009). Moreover, the use of AI may erroneously lead to price discrimination based on race or gender,

which is both illegal and immoral (Tam, 2018). Although we agree with these arguments, we believe that these issues are unlikely to hold for many industries. Most industries that use personalized pricing are likely to exist in the absence of personalized pricing and are also likely to adhere to the law. That being said, we do acknowledge that our two behavioral disutility components may be context specific. Regarding unfairness perception, consumers might view dynamic pricing for perishable services—such as a last-minute discount on an otherwise unoccupied hotel room—to be fairer than a discount to encourage customer loyalty. Judgments about utility may also depend on the stakeholders. When payments for societally supported common goods (e.g., health care) are based on affordability (analogous to progressive taxes), with unequal payments equivalent to pro-social subsidies, people may discount the disutility from perceived unfairness. Regarding surveillance aversion, people may weigh it by information sensitivity, e.g., more averse if based on recent drug prescription purchases than gender or occupation.

There is perhaps a more philosophical discussion appropriate in terms of what types of utility should be incorporated when we consider questions of ethics. Coker and Izaret (2021) chose not to incorporate disutility from either unfairness perception or surveillance aversion and instead took a more pragmatic view that utility from the product is enshrined in its objective use. We would argue, instead, that many non-objective traits of products perform a key role in reflecting utility to consumers. These include, for example, whether a product is sustainably sourced, or one experiences sentimental value toward a product or likes a product's ambassador. We would argue unfairness perception and surveillance aversion are similarly enshrined in a product and/or brand, and should thus be incorporated; however, we accept that this is a matter of perspective. That being said, our sensitivity analyses of the original Coker and Izaret (2021) assumptions showed that there are situations where we do not need any form of disutility from unfairness perception or surveillance aversion for welfare to be negatively affected by personalized pricing. In any case, we acknowledge that the decision of what types of utility should be incorporated in the utility function may be dependent on the time period, context, or culture.

Transparency of Price Personalization

Although survey evidence suggests that consumers indeed incur disutility from both our behavioral components, disutility is only experienced when the consumer finds out they have been monitored and/or received personalized pricing. In other words, our findings hinge on knowledge. We would argue that in the digital age, such information will invariably reach consumers and producers will not be able to hide price personalization, but we note such knowledge is a necessary element underlying our claims. Thus, the inclusion of our

behavioral components is restricted to those industries where price comparisons across consumers are possible, which would then allow for insights into whether price discrimination is implemented.

Extent of Price Personalization

In effect, the exposition provided by both Coker and Izaret (2021) and the current paper assume a world of quasi-personalized pricing. If the producer was a monopolist that was truly able to determine the WTP of both Alice and Bob, prices would be set in such a manner that producer welfare would be optimized—i.e., setting the price exactly to the WTP—which is not the case here. Namely, both Bob and Alice would still be willing to purchase the product for considerably higher unit prices than under the Coker and Izaret (2021) setup. A related view from the perspective of the monopolist under fully personalized pricing is provided in e.g., Bergemann et al. (2015). We argue it is reasonable that we will likely remain in such a quasi-personalized world as true WTP will be challenging to compute by AI algorithms—i.e., there will likely remain some residual noise in WTP predictions on the individual level. Given that there are typically many vendors in a market as well as regulations in place regarding market competition will also make fully personalized pricing unlikely in a non-monopolist setting. The dynamics of personalized pricing in the context of multiple vendors is an intriguing question which is beyond the scope of this paper as we simply assume that some form of personalized pricing has been realized. For example, it may be the case that services like Honey or Capital One—which allow consumers to compare prices—will enable consumers to sidestep personalized pricing. For our current purposes, we believe that such mechanisms will maintain the status quo of quasi-personalized pricing within which our analysis is based.

Discussion and Concluding Remarks

Assessing the ethicality of unitary versus personalized pricing under four consequentialist SWFs, we found that whether price personalization ethically outperforms unitary pricing depends on the relative value of equality and total utility under different SWFs. These findings contrast with earlier work by e.g., Coker and Izaret (2021) who found that personalized pricing improved social welfare regardless of which social welfare function one holds to be most plausible. Given their findings that personalized pricing increases both utility and equity, they concluded that personalized pricing is ethically superior to unitary pricing. We caution that their findings do not strictly hold when relaxing their assumption that the richer individual also obtains the higher utility

from the product—i.e., even without including disutility due to unfairness perception and/or surveillance aversion, their findings need not hold. On the contrary, it appears that unitary pricing is preferred in most cases if consumers perceive price personalization as unfair or they feel their privacy has been breached by the AI algorithms used to approximate WTP. This means that the finding of improved welfare due to personalized pricing as found by Coker and Izaret (2021) fundamentally depends on whether inequalities in utility from a product and disposable income align—i.e., that Alice, who is richer than Bob, also obtains more utility from the product. In other words, if algorithms personalize based on either the expected utility from a product or one's disposable income, it can happen that even without any disutility from unfairness perception or surveillance aversion, personalized pricing leads to a welfare loss. These findings are especially relevant given that product utility might very well increase for low-income consumers. More generally, our findings emphasize the importance of considering varied setups when tackling questions of ethicality as unexpected reversals of results may easily occur.

Using a four-components WTP decomposition, we further demonstrated that personalized pricing can be ethically outperformed by unitary pricing for small levels of disutility from unfairness perception and surveillance aversion under SWFs that tend to prioritize total utility (utilitarianism and prioritarianism). Personalized pricing ethically outperforms unitary prices under SWFs that tend to prioritize equality (egalitarianism and leximin), be it at the cost of overall utility. We also showed that the preference for personalized pricing under the leximin SWF strongly depends on the size of the shared disutility component (in our case surveillance aversion). Specifically, in the setup of Coker and Izaret (2021), we calculated that increasing surveillance aversion from 5% of the utility gained from the purchase to one-eighth is already enough to favor unitary pricing over personalized pricing. Therefore, we conclude that only in a rather extreme case of egalitarianism in which equality is prioritized above all else, personalized pricing consistently ethically outperforms unitary pricing.

Although we provide insights across a range of parameter settings, we identify several avenues for further research. First, our analysis could be extended by expanding the number of consumers in terms of income and personal utility to better reflect the complexity of an actual market. Our online tool allows the reader to vary our assumptions further, but analyses of large-scale population of strongly heterogeneous consumers remain as an extension. We do note that our results are only dependent on large-scale population dynamics insofar as these might affect the disutility sources incorporated in our decomposition. Relatedly, incorporating heterogeneity and population dynamics in terms of the disutility

due to unfairness perception and surveillance aversion is an interesting avenue of future research (e.g., unfairness disutility becoming conditional on one's relative position in terms of differential pricing etc.). Second, our analysis includes simplifying assumptions on the part of the producer. Additional market dynamics could be introduced by expanding the number of firms in the market and including optimization of price-setting on the part of the producer. Third, we have used estimates from lab experiments to identify reasonable estimates for disutility due to unfairness perception and have set similar estimates for surveillance aversion. However, these estimates are likely to differ across markets, both in terms of the product being sold, as well as possible cultural differences across markets (Bolton et al., 2010). As mentioned above, similar heterogeneity might be present at the level of the consumer. Providing additional empirical estimates of these sources of disutility and their variability across contexts and time periods is a third avenue of research that should further inform our understanding of the ethicality of personalized pricing.

Finally, it is useful to reflect on how our simulation analysis relates to the real world. The goal of our simulation analysis was to show that the previous conclusion that personalized pricing ethically outperforms unitary pricing was not uniformly true even under the original Coker and Izaret (2021) assumptions, and even more so when including the two additional disutility components. As such, we manipulated each of the original parameters in turn. However, in a real-world setting, it is plausible to assume that all parameters should be changed simultaneously. Thus, although we do provide initial insights into the dynamics and interplay of these parameters, future research should go even further in mimicking a market as it is in the real world.

Already in their current form, our findings illustrate that increasingly negative perspectives toward personalized pricing may be warranted. The potentially higher profits resulting from personalized pricing for the producers may be undermined by the lower social welfare for the consumers, whenever they incur disutility from the process. We already show this in the stylized case of equal producer welfare and comparatively small levels of disutility on the part of the consumer. Therefore, policy discussions should take these factors into consideration when considering whether personalized pricing should be limited or adjusted in the future. As such, this paper serves as a call for a reconsideration of the premise that personalized pricing is always ethical.

Appendix

See appendix Figs. 6, 7




Fig. 6 Welfare under the Coker and Izaret (2021) setup without disutility from unfairness perception and surveillance aversion when varying both personal utility of and unit price charged to player A. Each

column further represents different starting values for the disposable income of both players (Coker and Izaret (2021) income assumption in the middle column). Each row reflects a different welfare function




Fig. 7 Welfare under the Coker and Izaret (2021) setup without disutility from unfairness perception and surveillance aversion when varying both personal utility of and unit price charged to player B. Each

column further represents different starting values for the disposable income of both players (Coker and Izaret (2021) income assumption in the middle column). Each row reflects a different welfare function

Declarations

Conflict of interest The authors have not disclosed any competing interests.

Research Involving Human and Animal Participants Not applicable.

Informed Consent Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Acquisti, A., & Grossklags, J. (2008). 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). Auerbach Publications.
- Allender, W. J., Liaukonyte, J., Nasser, S., & Richards, T. J. (2021). Price fairness and strategic obfuscation. *Marketing Science*, 40(1), 122–146.
- Anderson, E. T., & Simester, D. I. (2010). Price Stickiness and customer antagonism. *The Quarterly Journal of Economics*, 125, 729–765.
- Arrow, K. (1951). *Social choice and individual values*. New York City: John Wiley & Sons.
- Bar-Gill, O. (2019). Algorithmic price discrimination when demand is a function of both preferences and (Mis)perceptions. *The University of Chicago Law Review*, 86(2), 217–254.
- Bentham, J. (1780). *An Introduction to the Principles of Morals and Legislation*.
- Bergemann, D., Brooks, B., & Morris, S. (2015). The Limits of price discrimination. *American Economic Review*, 105(3), 921–957.
- Bergson, A. (1938). A reformulation of certain aspects of welfare economics. *Quarterly Journal of Economics*, 52(2), 310–334.
- Bolton, L. E., Keh, H. T., & Alba, J. W. (2010). How do price fairness perceptions differ across culture? *Journal of Marketing Research*, 47(3), 564–576.
- Coker, J., & Izaret, J.-M. (2021). Progressive pricing: The ethical case for price personalization. *Journal of Business Ethics*, 173, 387–398.
- Culnan, M. J., & Bies, R. J. (2003). Consumer privacy: Balancing economic and justice considerations. *Journal of Social Issues*, 59(2), 323–342.
- DellaVigna, S., & Malmendier, U. (2006). Paying not to go to the gym. *American Economic Review*, 96(3), 694–719.
- Elegido, J. M. (2011). The ethics of price discrimination. *Business Ethics Quarterly*, 21(4), 633–660.
- Hermann, E. (2022). Leveraging artificial intelligence in marketing for social good—an ethical perspective. *Journal of Business Ethics*, 179, 43–61. <https://doi.org/10.1007/s10551-021-04843-y>
- Kahneman, D., Knetsch, J. L., & Thaler, R. (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review*, 76(4), 728–741.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–292.
- Krämer, A., Friesen, M., & Shelton, T. (2018). Are airline passengers ready for personalized dynamic pricing? *Journal of Revenue and Pricing Management*, 17, 115–120.
- Li, K. J., & Jain, S. (2016). Behavior-based pricing: An analysis of the impact of peer-induced fairness. *Management Science*, 62(9), 2705–2721.
- Loi, M., Hauser, C., & Christen, M. (2022). Highway to (Digital) Surveillance: When are clients coerced to share their data with insurers? *Journal of Business Ethics*, 175, 7–19. <https://doi.org/10.1007/s10551-020-04668-1>
- Lukosius, V., & Hyman, M. R. (2018). Personal Internet Shopping Agent (PISA): A Framework.
- Marcoux, A. (2006). Much ado about price discrimination. *Journal of Markets and Morality*, 9(1), 57–69.
- Martin, N. (2019, March 30). *Uber Charges More If They Think You're Willing To Pay More*. Retrieved from Forbes: <https://www.forbes.com/sites/nicolemartin1/2019/03/30/uber-charges-more-if-they-think-youre-willing-to-pay-more/?sh=9fc7caa73654>
- Martin, K. (2016). Understanding privacy online: development of a social contract approach to privacy. *Journal of Business Ethics*, 137, 551–569.
- Mattioli, D. (2012, August 23). *On Orbitz, Mac Users Steered to Pricier Hotels*. Retrieved from The Wall Street Journal: <https://www.wsj.com/articles/SB10001424052702304458604577488822667325882>
- Mikians, J., Gyarmati, L., Erramilli, V., & Laoutaris, N. (2012). Detecting Price and Search Discrimination on the Internet. *Proceedings of the 11th ACM Workshop on Hot Topics in Networks* (pp. 79–84). New York City: Association for Computing Machinery.
- Miller, A. A. (2014). 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.
- Miller, K. M., Hofstetter, R., Krohmer, H., & Zhang, Z. J. (2011). How should consumers' willingness to pay be measured? An empirical comparison of state-of-the-art approaches. *Journal of Marketing Research*, 48(1), 172–184.
- Parfit, D. (1998). Equality and priority. In A. Mason (Ed.), *Ideals of equality* (pp. 1–20). Basil Blackwell.
- Priester, A., Robbert, T., & Roth, S. (2020). A special price just for you: Effects of personalized dynamic pricing on consumer fairness perceptions. *Journal of Revenue and Pricing Management*, 19, 99–112.
- Rawls, J. (1971). *A theory of justice*. Harvard University Press.
- Rothschild, P., Boudet, J., & BenMark, G. (2019). Why Personalization Matters for Consumer Privacy. *MIT Sloan Management REVIEW*.
- Seele, P., Dierksmeier, C., Hofstetter, R., & Schultz, M. D. (2020). Mapping the ethicality of algorithmic pricing: A review of dynamic and personalized pricing. *Journal of Business Ethics*, 170, 1–22.
- Sen, A. (1979). Equality of what? *The tanner lectures on human values*.
- Shead, S. (2017, June 5). *Uber has refunded customers caught up in London terror attack*. Retrieved from Business Insider: <https://www.businessinsider.com/uber-refunds-customers-surge-pricing-london-terror-attack-2017-6#:~:text=Reuters%20LONDON%20%E2%80%94%20Uber%20has%20refunded,criticised%20for%20charging%20surge%20fares.&text=%22We%20are%20ensuring%20all%20rides,in%20Londo>

- Shiller, BR (2016). Personalized price discrimination using big data. *Brandeis university working paper series*, 108, 1–39.
- Snyder, J. (2009). What's the matter with price gouging? *Business Ethics Quarterly*, 19(2), 275–293.
- Steinberg, E. (2020). Big data and personalized pricing. *Business Ethics Quarterly*, 30(1), 97–117.
- Tam, L. (2018, February 26). *Price discrimination based on race is ugliest - just ask the Chinese at Heathrow Airport*. Retrieved from South China Morning Post: <https://www.scmp.com/news/hong-kong/community/article/2134704/price-discrimination-everywhere-ugliest-when-race-involved>
- Temkin, L. (2000). *Equality, Priority, and the Levelling Down Objection*.
- Turow, J., Feldman, L., & Meltzer, K. (2005). Open to Exploitation: America's Shoppers Online and Offline. *A Report from the Annenberg Public Policy Center of the University of Pennsylvania*, Retrieved from http://repository.upenn.edu/asc_papers/35.
- Turow, J., King, J., Hoofnagle, C. J., Bleakley, A., & Hennessy, M. (2009). *Contrary to what marketers say, Americans Reject Tailored Advertising and Three Activities That Enable It*. The New York Times. Retrieved from https://graphics8.nytimes.com/packages/pdf/business/20090929-Tailored_Advertising.pdf
- Turow, J., Hennessy, M., & Draper, N. (2015). How Marketers Are Misrepresenting American Consumers And Opening Them Up to Exploitation. *A Report from the Annenberg Public Policy Center of the University of Pennsylvania*, Retrieved from file:///C:/Users/deni/Desktop/SSRN-id2820060.pdf.
- Varian, H. R. (1992). *Microeconomic analysis*. W.W. Norton.
- Xia, L., Monroe, K. B., & Cox, J. L. (2004). The price is unfair! A conceptual framework of price fairness perceptions. *Journal of Marketing*, 68(4), 1–15. <https://doi.org/10.1509/jmkg.68.4.1.42733>
- Zuiderveen Borgesius, F., & Poort, J. (2017). Online price discrimination and eu data privacy law. *Journal of Consumer Policy*, 40, 347–366.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.