



Progressive Pricing: The Ethical Case for Price Personalization

Jerod Coker¹ · Jean-Manuel Izaret²

Received: 5 September 2019 / Accepted: 28 May 2020 / Published online: 4 June 2020
© The Author(s) 2020, corrected publication July 2020

Abstract

Price discrimination is widely considered unethical/unfair by consumers, as has been borne out by decades of psychological research and mainstream press reporting. However, little academic work has been done to investigate the ethics of price discrimination. The work that has been done to date concludes that while price discrimination is not unethical, despite widespread lay perceptions, it is at best morally neutral. We argue price discrimination is *more* ethical than unitary pricing, when done ‘progressively,’ meaning firms charge customers as a function of their willingness-to-pay. We introduce this specific kind of price discrimination as ‘Progressive Pricing’ and demonstrate it ethically outperforms a ‘Unitary Pricing’ scheme (where everyone pays the same price, regardless of their willingness-to-pay), at least within a broadly consequentialist framework. We do this by comparing a Unitary Pricing scheme to a Progressive one, analyzing them through the lenses of four different consequentialist ‘Social Welfare Functions’ (Utilitarian, Egalitarian, Prioritarian, and Leximin), which are used by welfare economists and philosophers to rank the distributions of different social outcomes, concluding that Progressive Pricing is preferred regardless of which Social Welfare Function(s) one finds most plausible.

Keywords Price discrimination · Progressive Pricing · Leximin · Egalitarianism · Prioritarianism · Utilitarianism · Social welfare function · Willingness-to-pay · Price personalization · Dynamic pricing · Consequentialism

Introduction

The explosion of big data, increasing sophistication of machine-learning tools, and the growing popularity of online shopping have created the conditions for ‘price personalization’ by online merchants, in which individual customers are given different prices for the same products based on their specific consumption behaviors and preferences (Ayadi et al. 2017). In industries where the marginal cost to production is trivial (as in music and video streaming services, the distribution of software, and social network applications), the pressure for such personalized pricing is much greater, as it

is impossible for firms to be profitable when prices are set near marginal costs (Phlips 1983).

This pricing is closely related to what economists classically refer to as ‘first degree’ or ‘perfect’ price discrimination (Shapiro and Varian 1999). In general, the idea of perfect price discrimination is that each consumer’s maximum willingness-to-pay (WTP) for a product or service is known, and that the supplier charges each consumer their maximum WTP, thereby extracting all surplus for themselves. In traditional welfare economic analysis, the results of perfect price discrimination are ambiguous—it is possible to set up situations in which total utility can be higher, lower, or the same under perfect price discrimination compared to other pricing regimes (Schmalensee 1981; Varian 1985). However, usually it is assumed that the supplier extracts all of the consumer surplus by setting their price exactly to maximum WTP (or as closely as possible), and this is generally perceived to harm consumer welfare, even while the *overall* welfare effects may be ambiguous.

The closest real-world approximations will differ from this textbook description, however. First, in reality, a customer’s maximum WTP will be proxied algorithmically and will have considerable margin of error. Second, WTP may change over time for each individual, depending on

The original version of this article was revised due to a retrospective Open Access order.

✉ Jerod Coker
Coker.Jerod@bcg.com

Jean-Manuel Izaret
Izaret.JeanManuel@bcg.com

¹ Boston Consulting Group, 515 South Flower Street, Suite 5100, Los Angeles, CA 90071, USA

² Boston Consulting Group, Two Embarcadero Center, Suite 2400, San Francisco, CA 94111, USA

their specific pattern of purchasing behavior. Finally, personalized prices need not be equal to the estimated maximum WTP—indeed, we will argue they should not be. We will therefore argue that price personalization need not be detrimental to consumer welfare if done in a way we call ‘Progressive Pricing,’ meaning producers as well as consumers can benefit. This is partly because, as we will show, firms should not charge all of a customer’s maximum WTP, not only because of the aforementioned uncertainty, but also because of a likely negative consumer response to having all their surplus extracted. The latter is a particularly salient concern given the public backlash to early discriminatory pricing scandals, such as the discovery in 2000 that Amazon was charging different prices to different consumers for the same DVDs, or that Staples in 2012 was charging customers different prices based on their zip code. Writing in the *New York Times*, Nobel Prize-winning economist Paul Krugman wrote in response to the Amazon controversy: “‘[D]ynamic pricing’ is about to become a major consumer issue.... The only thing that is likely to stop it is government action” (Krugman 2000). In the same op-ed, he makes a striking claim: “...[D]ynamic pricing is also undeniably unfair: some people pay more just because of who they are.”

Consumers generally agree: in survey after survey, respondents indicate strong aversion to paying different (especially higher) prices than other consumers for the same good (Turow et al. 2005). 76% of people surveyed, for example, agree with this statement: “It would bother me to learn that other people pay less than I do for the same products.” 87% disagreed with this statement: “It’s OK if an online store I use charges people different prices for the same products during the same hour.” While psychological research has concluded that consumers react less negatively when they are *advantaged* by the price discrimination (Ordonez et al. 2000), still 72% disagreed with the following: “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.”

We will argue that such intuitions are misguided, and that there is a way to implement price personalization in a way that is societally beneficial, at least from a consequentialist point of view. The basic premise of Progressive Pricing is that, rather than set a price equal to a buyer’s estimated maximum WTP, firms instead price as some *function* of that estimated maximum WTP (e.g., 50% of maximum WTP), hence changing the mindset from ‘value extraction’ to ‘value sharing.’ We will argue, if done in this Progressive way, there could co-exist personalized pricing along with higher and more equal consumer welfare. We will argue a Progressive Pricing scheme is more socially beneficial than the ‘Unitary Pricing’ (where everyone pays the same price) regime that most regard as intuitively fairer.

Ethics of Price Discrimination Background

Much research has been performed concerning consumer perceptions of the fairness of pricing and price discrimination. By and large, consumers find price discrimination unfair, even when it advantages them (Martins 1995; Huppertz et al. 1978). For academics, Moriarty (2017) summarizes, “[W]hile economists tend to think that price discrimination is valuable insofar as it enables firms to increase output...the moral status of it is less clear.” Indeed, unlike the research on consumer perceptions, very little has been written about the actual *ethics* of price discrimination. As Elegido (2011) writes, “[E]thicists, with very few exceptions (Marcoux 2006), have done almost no systematic work” on the topic of price discrimination. Ultimately, he and Marcoux both conclude that price discrimination is not per se unethical or unfair.

Marcoux’s basic argument is that an equal price actually affords *unequal* welfare to consumers with different reservation prices (i.e., maximum WTP), and so if ‘equal treatment’ is the barometer of ethicality, then price discrimination more plausibly serves that end than Unitary Pricing. That is, if one person has a maximum WTP of \$10 and another \$5 for the same good, and the unitary price for both is \$5, one has a relatively large surplus of \$5 (50%), whereas the other has no surplus whatsoever. However, he does not endorse the view that price discrimination is better than Unitary Pricing, but rather that *if* we believe equalizing welfare is what an equal treatment norm demands, and if we believe such an equal treatment norm exists for pricing, then it ought to be preferred. Additionally, he argues that a *procedural* perspective inclines neither for nor against price discrimination, since “Charge everyone their reservation price,” which requires price discrimination insofar as people have different reservation prices, is equivalently procedurally fair to “Charge everyone the same price.” In sum, he writes: “...the widely held view about the unfairness of price discrimination is untenable; fairness considerations incline either *no more against or else strongly in favor* of price discrimination, as against other pricing regimes” (p. 58) [emphasis added].

Elegido’s conclusion with respect to price discrimination is similar to Marcoux’s: “[I]n itself, price discrimination is a *morally neutral* practice that businesspeople are entitled to use if it advances their morally legitimate interests” [emphasis added]. He demonstrates that a blanket ban on price discrimination leads to implausible restrictions, for example when it is required for a business to exist or be profitable. His summarized view is that, “ultimately, there is no independent ethics of price discrimination... Whether or not there is price discrimination is, in itself, simply irrelevant to the justice of a price” (p. 654).

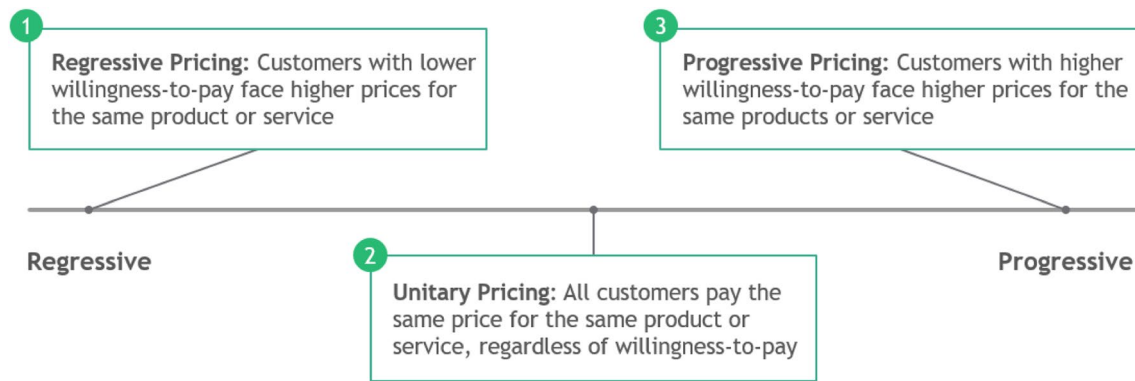


Fig. 1 The pricing progressivity scale

Rather, his view is that a price can be considered just if it is one that prevails in an open market (more on this in Sect. 6.2)—and so, if price discrimination occurs in a way where a ‘substantial portion’ of consumers pay a given price, it is fine, even if there are multiple prices paid.

We will build upon the pioneering work of Marcoux and Elegido, going a step further than their conclusion that price discrimination is morally neutral. Instead, we will argue that, when done ‘progressively’ (which is the natural way, i.e., where those with higher WTP pay higher prices), price personalization is, from a broadly consequentialist point of view, actually *better* for society overall than Unitary Pricing regimes.

Willingness-to-Pay (WTP) Decomposition

For this paper, we decompose willingness-to-pay (WTP) into two distinct components: the utility from money lost from a purchase and the utility gained from the purchase. What sets a rational person’s maximum WTP is therefore when the utility *lost* from the price, $u(L)$, is equal to the utility *gained* from the purchase $u(G)$. Of course, a rational individual would buy at any point where $u(G) > u(L)$ and would never buy when $u(G) < u(L)$.

This decomposition implies the maximum WTP for Person A can be higher than the maximum WTP for Person B for the same good or service at the same time for two reasons: A has a lower utility loss from an equivalent price paid (we will generally assume this means they are richer, due to an assumption of diminishing marginal utility of income and wealth) and/or A derives higher utility from the good or service. While there is debate as to whether and how much the marginal utility of income/wealth declines as it increases (Layard et al. 2008; Oswald 2005), we take the common view that the relationship between income and utility is that the utility derived from marginal increases in income

is ever-increasing, but that it does so at an ever-decreasing rate (e.g., the relationship can be described as logarithmic).

Of course, these two sides of the WTP coin need not go together, and when they conflict the results may net out in either direction. In other words, if Person C has higher income but lower utility from Good X compared to Person D, it is unclear whether their maximum WTP will be higher or lower. While economists typically discuss consumer surplus in terms of dollars (e.g., if a consumer’s max WTP is \$100 and their price is \$80, they are said to have \$20 in surplus), the ‘currency’ we will use is utility itself (Cohen 1989). This is because just knowing a consumer’s maximum WTP and their price will not be sufficient to compare *relevant* consumer surpluses, as \$20 could entail a very different level of utility for different individuals.

Progressive Pricing in Detail

Progressive Pricing can be thought of on a continuum of progressivity from ‘Regressive Pricing’ (whereby customers with lower WTP pay higher prices), ‘Unitary Pricing’ (the dominant view of ‘fair pricing’ by consumers, whereby all consumers pay the same price, regardless of WTP), and of course ‘Progressive Pricing’ (whereby price increases as WTP increases) (see Fig. 1).

An example helps illustrate each point on the scale. Imagine that Alice (a wealthy businesswoman) and Bob (an undergraduate business student) each subscribe to *The Business Journal*. Let’s assume their utility function for income is the same, and can be represented by the function $utility = \ln(income)$. (We choose the natural logarithm function as it satisfies the ‘diminishing marginal utility’ requirement well; that is, it has a constantly positive first derivative, but a constantly negative second derivative. However, other functions would work as well, e.g., the square root function.) Let us now presume Alice has an average income of \$25,000 per month. Using our utility function, then, the monthly utility

from her income is ~ 10.1 (let's imagine utility is on a scale from 0 to 100, with 100 being a perfectly content life and 0 being indifference between life and death). Notice that even at a very high income of \$25,000 per month, this only provides Alice with $\sim 10\%$ of her maximum possible utility, the rest of which would be derived from having good health, being safe, feeling self-actualized, etc. Bob, on the other hand, has an income of \$2500 per month for a monthly utility of ~ 7.8 from income.

Now, let us presume that *The Business Journal* gave Alice a corporate discount for a price of \$20 per month, while charging Bob the normal rate of \$30 per month. Let us further presume that Alice gets 0.1 utils from her subscription versus Bob's 0.05 utils; this is because Alice requires the *Journal* to do her job, whereas Bob only reads it to casually stay up to date with the goings on of the private sector.

Now, we can calculate the utility surplus of both Alice and Bob to compare. The utility lost for Alice from spending \$20 per month is $\ln(25,000) - \ln(24,980) = 0.0008$. This means her utility surplus is $0.1 - 0.0008 = 0.0992$. On the other hand, Bob loses $\ln(2500) - \ln(2470) = 0.0121$, netting $0.05 - 0.0121 = 0.0379$. Note the net utility surplus for Alice is ~ 2.6 times as great as Bob's, despite the fact she is already better off. Additionally, Alice pays only 0.8% of her utility gained (0.008 divided by 0.1), whereas Bob pays 16% of his (0.0121 divided by 0.05). This is clearly an example of Regressive Pricing, then, because Alice has a much higher maximum WTP, but a lower price, meaning both her absolute and proportional utility surpluses are significantly higher. This is not just because she has more income; if Alice and Bob had the same income, the fact that she gets double the utility from *The Business Journal* would also imply that if she paid less, the pricing would be regressive.

It should be intuitively obvious that the Regressive Pricing in this example is not optimal from a social welfare standpoint, just simply from understanding that a wealthy businesswoman pays less than a student who has much lower income and who does not derive as much utility from the subscription. This intuition is backed up by the reasoning above. However, we will soon show this *same* reasoning means that Unitary Pricing, the 'common sense' fair pricing regime, can result in large differences in consumer utility surplus as well. In other words, we will show that Unitary Pricing is worse than Progressive Pricing in the same way that Regressive Pricing is worse than Unitary Pricing (and, by transitivity, Progressive Pricing).

This is because a Progressive Pricing regime in which lower-WTP individuals pay less than higher-WTP individuals at least *narrows* this utility gap. It may be infeasible in practice to eliminate the gap completely (for example, even under the assumption that Alice and Bob both gain 0.1 utils from the subscription, Alice's price would need to be orders of magnitude higher than Bob's to result in equal welfare

diminution, based on Alice's much higher income alone). While Marcoux (2006) does not endorse this conclusion, he similarly mentions that "...equal welfare diminution...is not the case under a unitary pricing regime," continuing, "[I]f fairness demands that each buyer enjoy the same welfare from purchasing the same product, then some form of price discrimination (whether one that charges each his reservation price or another price that affords the same degree of consumer surplus) is necessary to achieve fairness." We will build upon this correct conclusion, arguing that Progressive Pricing is superior from a social welfare standpoint.


Progressive Pricing and Social Welfare Functions (SWFs)

While the intuitive appeal of Progressive Pricing should be apparent from the prior example, we now turn to comparing how this utility surplus comparison can be analyzed using a variety of 'Social Welfare Functions' (SWFs), which are functions that rank social states as less, more, or equivalently desirable; they are commonly used in welfare economic analyses. Indeed, welfare economists classically view normative economic analysis from both 'efficiency' and 'equity' lenses, generally with the assumption that both efficiency (usually meaning total utility) and equity are socially desirable goals, but that often there is a tradeoff between the two normative ideals—e.g., most policies will increase total utility at the expense of equality or vice versa. However, there is debate as to how much this 'Efficiency-Equity Tradeoff' manifests in reality (Sachs 2017), and indeed we will argue that Progressive Pricing is both more efficient *and* equitable, meaning it is socially desirable whether one is a strict utilitarian, egalitarian, or like many normative economic analyses, somewhere in the middle.

From a welfare economics and philosophy of economics point of view, there are a few plausible and widely discussed SWFs, all of which are fundamentally consequentialist in nature. In this paper, we will not delve into issues of meta-ethics by arguing that one ought to favor a broadly consequentialist framework over any of the competing ethical paradigms—we will instead concede that if one were to reject consequentialism broadly, then the arguments presented will not be compelling. However, importantly, we do not require readers to choose a framework *within* consequentialism (e.g., utilitarianism); instead, we will introduce each of the aforementioned four social welfare functions briefly and demonstrate how Progressive Pricing outperforms Unitary Pricing under each, concluding that Progressive Pricing outperforms Unitary Pricing *regardless* of which plausible SWF one prefers.

For each SWF, we will compare how a pricing of \$20 each compares to a pricing of \$35 for Alice and \$5 for Bob.

Fig. 2 Demand curve for a high fixed cost, low marginal cost software company



We do this to hold the *Journal's* own supplier welfare constant (assuming \$40 total is the same utility for it, regardless of from whom it is paid), as we are primarily interested in exploring *consumer* utility surplus in this paper.

Utilitarianism

Many people are familiar with the Utilitarian SWF. Its goal is to add up the total utilities of every person (or sentient creature) in the world (or universe) and to maximize the sum total of their utilities. For example, if we only consider the utility of Alice and Bob, and sticking with the 0 to 100 scale described previously, the closer their total utility is to 200, the better is the state of the world.

Progressive Pricing is clearly to be favored under a Utilitarian SWF in this case, as the sum total of utility under a Unitary Pricing scheme of \$20 each results in lower aggregate utility than charging Alice \$35 and charging Bob \$5, again assuming Alice's utility gained is 0.1 and Bob's is 0.05. To quickly calculate, the utility surplus distribution for (Alice, Bob) under Unitary Pricing of \$20 each is (0.0992, 0.0420),¹ the sum of which is 0.1412. Now compare this to the Progressive Pricing scheme, with pricing of \$35 and \$5 for Alice and Bob, respectively: the utility distribution for (Alice, Bob) becomes (0.0986, 0.0480), for a sum of 0.1466.² From a Utilitarian perspective, this is a ~4%

improvement, as the sum is higher in the latter scenario. Recall, this improvement is non-zero sum with the producer surplus, as in both cases the *Journal* receives \$40.

Of course, there is an added Utilitarian benefit to Progressive Pricing, beyond just considering the impact on Alice and Bob. Indeed, it can be shown that in real industries, some form of price discrimination is *required* for the market to function at all, and therefore for any consumers to achieve any utility surplus. This is true in many businesses with high fixed costs, but low marginal costs, such as streaming, software, and pharmaceutical companies. Philips (1983, p.1), for example, writes that, "[g]enerally, discriminatory prices [are] required for an optimal allocation of resources in real life situations." Similarly, Elegido references the economics research of Baumol and Swanson (2003) and Levine (2002), suggesting that price discrimination "may easily result in better outcomes for everyone." This, of course, satisfies a Utilitarian SWF, and indeed the dominant framework of welfare economics remains a Utilitarian one, hence the reason many welfare economists are not put off by price discrimination: if an industry or firm depends on price discrimination to exist, and that business's existence creates consumer and producer surplus (utility), then clearly the state of the world in which that business exists and price discriminates is better from a Utilitarian perspective than one in which it cannot price discriminate and therefore cannot exist. In the typical parlance of welfare economists, this is a 'Pareto Improvement,' as everyone is better off and no one is worse off.

For example, imagine a software business with \$100,000 in startup costs, and a marginal cost of \$1 per unit. If you review Fig. 2 below, you will see a contrived demand curve whereby there is demand for 20,000 units at a price of \$1; 10,000 units at a price of \$5; 5000 units at a price of \$10; and 2500 units at a price of \$30. It can be easily

¹ Alice's surplus is 0.1 utility gained— $[\log(25,000) - \log(24,980)] = 0.0992$; Bob's is 0.05 utility gained— $[\log(2500) - \log(2480)] = 0.0420$.

² Alice's surplus is 0.1 utility gained— $[\log(25,000) - \log(24,965)] = 0.0986$; Bob's is 0.05 utility gained— $[\log(2500) - \log(2495)] = 0.0480$.

shown that at one single price point, there is no way for the company to make a profit. Indeed, at a uniform price of \$1, there is clearly no margin to make up for the startup costs (as marginal cost is \$1). At a price of \$5, there is a per-unit gross margin of $\$4 = \$5 - \$1$, so gross profits are $\$4 \times 10,000 = \$40,000$, which is insufficient to cover the startup costs. The same is true at a price of \$10 (gross margin of \$45,000) and \$30 (gross margin of \$72,500). However, with price discrimination, the software provider is able to earn a small profit of \$15,000.

Therefore, Progressive Pricing is to be preferred to Unitary Pricing under a Utilitarian SWF as it has higher total utility surplus. It should especially be preferred if it is required for a particular firm to survive, which is increasingly the case as more and more new firms face high fixed and low variable costs.

Egalitarianism

Utilitarianism is, of course, far from universally accepted by philosophers or laypeople (Rawls 1971). For one, it fails to consider equality at all. As an example, imagine now that the utilities of Alice and Bob could be 50 and 50, respectively, or else 99 and 1 (on our 0 to 100 scale). Most would prefer the former, equal arrangement. However, Utilitarianism fails to give us a guide; it views them equally, as they both sum to 100. On the other hand, a strictly Egalitarian SWF prioritizes equality above all else, and will always choose the more equal of two arrangements, hence preferring the 50/50 scenario over the 99/1.

How does Progressive Pricing then fare under an Egalitarian SWF? Again, it outperforms Unitary Pricing. Indeed, insofar as we consider utility the correct currency of justice, a Progressive Pricing scheme results in a more balanced utility distribution than Unitary Pricing. For Alice and Bob, above, recall the Unitary Pricing distribution is (0.0992, 0.0420), whereas the Progressive Pricing distribution is (0.0986, 0.0480). While the surplus gap is not eliminated, it is narrowed from 0.0572 to 0.0506, an improvement of ~11.5%. Again, this improvement is at no cost to the supplier.

To generalize, when prices are higher for those with higher WTP and lower for those with lower WTP (what we have dubbed Progressive Pricing), utility surpluses are closer to equality than under Unitary Pricing, where some consumers have all of their utility extracted (i.e., those whose WTP is exactly the unitary price have no surplus at all), while those with the highest WTP have large utility surpluses. Therefore, Egalitarians ought to join Utilitarians in preferring Progressive Pricing over Unitary Pricing.

Prioritarianism

Like Utilitarianism, however, the kind of simplistic Egalitarianism above is not widely accepted by philosophers, for the simple fact that it does not take aggregate utility into consideration at all. For example, in a scenario of (40, 90) versus (40, 40), it prefers the latter. However, to most this is repugnant; why should those with higher utility be sacrificed merely to equalize outcomes, without making anyone better off? Since a move from (40, 40) to (40, 90) is a Pareto Improvement (no one is made worse off and at least someone is made better off), economists and philosophers alike are reticent to deny such a move. This is commonly known as the ‘Levelling Down Objection’ to Egalitarianism (Temkin 2000).

Fortunately, there are SWF options that value more equal distributions but also more total welfare. One prominent example is Prioritarianism (Parfit 1998). It is similar to the Utilitarian SWF in that it seeks to maximize an overall goal, however, it does not suffer from the indifference Utilitarianism does of weighting all utility equally. Rather, it ‘priority-weights’ (hence the name) utility based on how much utility an individual already has. In other words, one can think of the Prioritarian SWF as itself a function of utility, whereby ‘Total Moral Value’ (in other words, the amount of what matters) is, for example, the logarithm of utility. This means that it is always better for there to be more utility, but moving from a 10 to a 20 is much more valuable than a 90 to 100, say (see Fig. 3).

Again, Progressive Pricing fares better than Unitary Pricing under a Prioritarian SWF. Indeed, it fares even better, *ceteris paribus*, under a Prioritarian SWF compared to a Utilitarian one, as the ‘moral value’ is now in a sense logged twice—that is, the diminishing marginal utility of income, combined with the diminishing marginal ‘moral value’ of utility, means that Alice’s welfare gains are quite heavily taxed. To illustrate, the Prioritarian distribution of increased ‘Total Moral Value’ under Unitary Pricing³ for (Alice, Bob) is (0.0097, 0.0054), for a sum of 0.0151; compare this to the Progressive Pricing⁴ distribution of (0.0097, 0.0061), for a

³ Unitary Pricing: Utility before for Alice was $\log(25,000) = 10.1266$. After her net utility increases from her subscription, this increases to 10.2258. This priority-weighted change is $\log(10.2258) - \log(10.1266) = 0.00975$. For Bob, utility before is $\log(2500) = 7.8240$. After his net utility increases from his subscription, this increases to 7.8660, for a priority-weighted change of $\log(7.8660) - \log(7.8240) = 0.0054$.

⁴ Progressive Pricing: Utility before for Alice is the same, 10.1266, as is Bob’s 7.8240. With the \$35 price, Alice’s utility after the subscription decreases slightly to 10.2252; with the \$5 price, Bob’s utility after the subscription increases to 7.8720. The priority-weighted change for Alice is $\log(10.2252) - \log(10.1266) = 0.00969$; Bob’s is $\log(7.8720) - \log(7.8240) = 0.0061$.

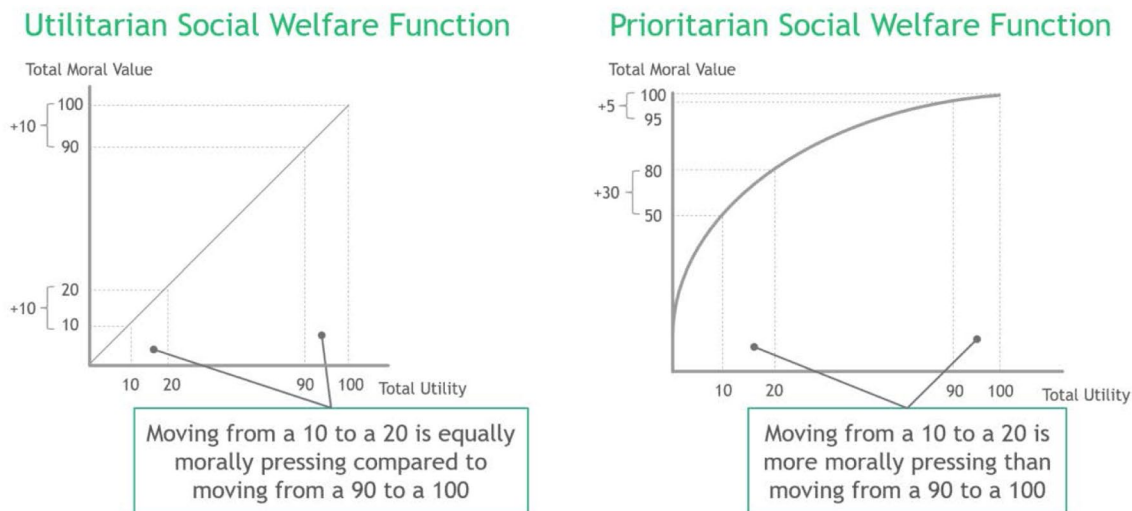


Fig. 3 The Utilitarian versus Prioritarian social welfare functions

sum of 0.0158: a ~5% improvement. Notice Alice’s priority-weighted utility barely decreases (by -5.8680×10^{-5}) due to the heavy taxing mentioned previously.

This Prioritarian outcome should not come as a surprise. Given that it is in a sense ‘between’ a Utilitarian SWF (its goal is to maximize a sum) and an Egalitarian SWF (it prefers when those who are worse off receive a given amount of utility), and given we have shown that both a Utilitarian and an Egalitarian SWF prefer Progressive Pricing over Unitary Pricing, it is natural that it should do the same.

Leximin

The final SWF we test Progressive Pricing against is derived from John Rawls’ (1971) ‘Maximin’ criterion, which states that the most desirable end-state is the one that maximizes the utility of the least well off individual. Amartya Sen (1980) built upon this idea to introduce the idea of ‘Leximin,’ which essentially adds a ‘tie-breaking’ condition whereby if the least well off is equally well off in two different scenarios, then the *next*-least well off takes lexical priority (hence the name Leximin). In other words, if the two states of the world could be (40, 40) or (40, 90), Leximin prefers the latter, because although the least well off is equally well off in both situations (40), the next least well off is better off in the latter (90 vs. 40).

Using this SWF, Progressive Pricing again outperforms Unitary Pricing. To continue our *Journal* example, recall that Alice and Bob have a Unitary Pricing net utility increase distribution of (0.0992, 0.0420), compared to the Progressive Pricing distribution of (0.0986, 0.0480). Notice that in Unitary Pricing, the least well off person (Bob) has a utility of 0.0420, whereas under Progressive Pricing, the least well off person (still Bob) has a utility of 0.0480—a ~14%

increase. Again, then, Progressive Pricing fares better under a Leximin (or Maximin) SWF, due to its effect on the least well off. This is intuitively obvious as well; in a Unitary Pricing scheme, some pay very close to their maximum WTP, and therefore achieve no consumer surplus. These are the least well off from the pricing scheme, clearly, while those who have very high WTP are the most well off under Unitary Pricing, as they have the most consumer surplus. Under Progressive Pricing, on the other hand, those with lower WTP pay less than they would under a Unitary Pricing regime (including, importantly, some that were priced out of the market completely, who would like the product or service but could not purchase it at the unitary price), clearly resulting in higher surplus for the least well off under Progressive Pricing.

Potential Objections

Price Discrimination is a Violation of an Equal Treatment Norm

We should not ignore the common ‘gut reaction’ that two people paying different prices for the same thing is unjust, which was described in the opening of this paper. Although this reaction is common among laypeople as previously described, and although some academics assert, “From an ethical point of view, charging different prices for the same product is a violation of the equal treatment norm that underlies market exchanges” (Ayadi et al. 2017), Elegido (2011, p. 641) correctly summarizes the state of ethical inquiry into a supposed equal treatment norm in pricing: “...no good arguments have been provided in the literature on equality to support the position that an equal treatment norm applies

to commercial transactions, and, more specifically, to pricing issues, and...there are important considerations that urge against such a norm in this context.” In other words, this common lay reaction has never been successfully defended as a rigorous philosophical principle.

Even if someone were to successfully argue that there is an equal treatment norm that applies to commercial transactions or pricing issues, it is not clear that such equal treatment would require equal pricing. Indeed, as we (and Marcoux (2006) before us) have argued, if one grants that the currency of egalitarian justice is utility (priority-weighted or otherwise), an equal treatment norm would *more* plausibly apply to utility (e.g., the rule ‘equalize utility’) than prices themselves (e.g., the rule ‘equalize prices’).

The Just Price as the Price in an Open Market

Elegido (2011, 2015) argues the case that a just price is that which is obtainable in an open market, by which he does *not* mean a competitive market, but rather one devoid of a legal or natural monopoly. His argument is roughly that when there are enough buyers to sustain a particular business at a particular price, then the prevailing market price (or a lower price) is a just one. Since the market can clearly bear that price, Elegido argues that is a strong sign of the inherent value of the product or service. He is careful to distinguish value from utility; in other words, while every individual may gain a different level of utility, there is one level of value, as demonstrated by the market price. Elegido (2011) gives us an example where his standard allows for price discrimination (p. 653):

You arrive in a city late at night and check into a hotel, at which time you are informed that the rate per night is \$300. You find it steep, but it is late, you are tired, and outside it is raining heavily, so you accept and move into your room. The following morning, while having breakfast, you strike up a casual conversation with a fellow guest and learn that she had made her reservation online and is only paying \$80 per night.

His argument is that if a ‘substantial portion’ of guests are paying the \$300 price, that is a good signal of its value and therefore it is a just price to charge, regardless of whether or not some other guests are paying a lower price. On the other hand, if some guest were to be charged, say, \$500, when the accepted open market price being paid by a substantial portion of guests is \$300, this would be morally suspect.

The issue we take with the ‘substantial portion’ criterion is that, under our Progressive Pricing model, it is possible and perhaps even likely that applying the same rule of, for example, ‘price equal to 50% of an individual’s WTP’ could easily lead to a continuum of prices on which no two individuals pay the same price (e.g., if there is complete

heterogeneity of individual WTP). Our view is of course that a ‘50% of WTP’ rule for pricing is superior to one in which there is one price or even a few discrete ‘buckets’ of consumers, with the same price within a bucket (what economists would call ‘third degree’ price discrimination). However, only in the latter scenario of third-degree price discrimination would a demand curve with highly heterogeneous WTP allow for price discrimination while also allowing for ‘substantial portions’ of consumers to pay the same price. Therefore, we conclude the ‘substantial portion’ rule to be overly-restricting: it could be the case that no two people pay the same price, and therefore there is no ‘substantial portion’ paying the price, and yet, within our framework, it could be perfectly acceptable.

In other words, this standard would not allow for Progressive Pricing as we have defined it. To illustrate, imagine a demand curve with 100 consumers, each with a distinct WTP ranging from 1 to 100, in one-unit increments. Imagine the pricing rule is ‘50% of WTP.’ The ‘substantial portion’ standard would not allow the consumer with WTP 100 to pay a price of 50, because no one else (and therefore no ‘substantial portion’ of consumers) would be paying such a high price. Therefore, their price would need to be reduced, let’s say to match the next-highest price, 49.5 (50% of the next-highest-WTP consumer’s WTP of 99). Although there is no explicit rule to identify what constitutes a ‘substantial portion,’ let’s presume 2 individuals each paying 49.5, when everyone else pays a lower price, is insufficient to meet that standard, and so those two must now have their prices reduced to 49 each, so now there are three consumers with the price 49, and so on until a ‘substantial portion’ is reached. We would submit this is an unnecessary reduction of price for the highest-WTP consumers; indeed, we have argued that a Progressive Pricing rule of ‘pay 50% of WTP,’ which would have the highest-WTP consumer pay 50 (more than any other consumer) is perfectly morally justifiable, regardless of the existence of other consumers paying such a high price or not. Indeed, with that rule, the highest-priced customers are also precisely the highest-*surplus* customers. In our example, the person with WTP of 100 pays 50, but keeps 50 surplus, meaning that, although she pays the highest price, she also achieves the highest (monetary) surplus.

Why is the Willingness-to-Pay Higher for Some?

Another potential objection to a Progressive Pricing regime is that it does not question why one individual’s utility is higher than another’s for a given good or service. We return to our WTP decomposition framework to tackle this objection in two parts. Recall that the two components of WTP are utility lost from the money spent (which we argue can be proxied by income: the higher one’s income, the lower the loss of utility from a given amount of money spent) and

the utility gained from the product or service. We begin with income.

If a person's WTP is higher because their income is higher, our argument is this should be presumed a just basis upon which to price discriminate. In other words, if two people derive the same level of utility from a product or service, but one has a higher income than the other, it need not be questioned that the person with the higher income pay a higher price, since there is a strong presumption they would lose less utility from an equivalent price paid and hence Unitary Pricing would result in differential surplus (which we have shown ought not be preferred under any plausible SWF). Of course, this does not mean that *any* higher price can be justified on the basis of higher income (using the prior reasoning, there is of course some higher price that will be *too* high so as to leave the higher income consumer with a lower utility surplus), but rather that the mere existence of some higher price for a consumer with higher income is not intrinsically suspect. However, we will argue that more scrutiny may be required when WTP is higher due to a higher utility derived from the product or service.

There is a long ethical literature and corresponding legal history concerning the immorality and illegality of price gouging, exploitation, and false advertising. For example, during a natural disaster there is often a surge in demand for water, gasoline, fuel, flashlights, etc., which would cause prices to rise in a free market. However, it should be clear that such price increases are exploitative and immoral (e.g., Snyder 2009a)—although there is some debate among philosophers (Zwolinski 2008; Snyder 2009b). A broader theory of exploitation (of which price gouging is arguably a form) has also been utilized to evaluate the wrongness of (some forms of) price discrimination, from St. Thomas Aquinas to Locke to today (Zwolinski and Wertheimer 2017). A famous example from Locke examines the case of two ships, one of which has lost all its anchors (and therefore cannot safely come to a stop or weather a storm). The ship comes upon another ship, which has an anchor to spare. Of course, the question becomes, at what price ought the captain of the extra-anchored ship offer her surplus anchor to the captain of the distressed ship? Locke's answer: "[T]he same price that she would sell the same anchor to a ship that was not in that distress. For that is still the market rate for which one would part with anything to anybody who was not in distress and absolute want of it" (1661/2004, p. 446).

On a related note, in discussing pricing based on a consumer's maximum WTP (as we have argued ought to be done), Elegido (2011, p. 640) states, "...it seems fundamentally wrong to make an argument about fairness in pricing depend on the buyer's reservation price. Along those lines, the more desperate the buyer happens to be, the higher the price the seller is justified in charging. In fact, perversely, it would follow that by not charging a very high price to a

customer who finds himself in a situation of very special need, the seller would be behaving unfairly."

We agree with the conclusions above, and in light of them argue it is not *always* the case that those with higher WTP should pay more, for example when such pricing would be an instance of price gouging specifically or exploitation more broadly. However, these considerations should be considered 'side constraints,' leaving the core thrust of our argument unscathed. In other words, it is not immoral to charge those with higher WTP higher prices in and of itself, but of course when such pricing would entail price gouging, exploitation, or any other condition that would make a price unjust between two individuals, one ought not implement it. However, these side constraints do the 'moral work,' not a prohibition on price personalization.

To summarize, why someone's WTP is higher than another's is indeed an important normative consideration. When it is because of a higher income, we argue it is not morally suspect. However, when it is because of a higher (expected) utility, one must investigate whether pricing in accordance with that higher (expected) utility would violate a moral side constraint such as avoiding price gouging, exploitation, or any other condition that would make the price of a transaction between two parties unjust. Our argument instead is that, given that the interparty dynamics of a transaction are just, Progressive Pricing increases social welfare, and as such it would be erroneous to limit its application absent these side constraints.

How do Firms Know About Differential Willingness-to-Pay?

Another set of 'side constraints' will likely apply to how firms come to estimate differential WTP for different customers, with some methods being clearly permissible (e.g., customers who willingly and knowingly sign up for loyalty programs) and others being clearly impermissible (e.g., installing a keystroke logger on their computer without their consent). In the digital age, we expect the impermissible methods of most concern will correspond to invasions of reasonable expectations of privacy on the Internet.

Indeed, in 2016 Europe passed the sweeping General Data Protection Regulation (GDPR); California followed in 2018 by passing the California Consumer Privacy Act, with more US states following thereafter. These regulations are designed to protect consumer data and privacy on the Internet, including the establishment of new rights such as 'data portability' (you must be allowed to export your personal data from one social media site or online retailer to another, e.g.) and 'the right to be forgotten' (they must delete your data if you request), among others. Given the recency of such laws, little academic work has been done on the practical implications of GDPR on price personalization,

although the scholarship to date concludes GDPR “generally applies to personalized pricing” and “[i]f enforced...could thereby play a significant role in mitigating any adverse effects of personalized pricing” (Zuiderveen Borgesius and Poort 2017). We will not take a stance in appraisal of these laws, but will instead argue that insofar as there are rights to privacy on the Internet, firms ought to respect those rights when algorithmically estimating WTP.

More work needs to be done to clearly define under which circumstances and how customer data can be collected, stored, and used ethically and legally, whether for use in pricing or otherwise. In the meantime, clearly there will be many instances in which a firm *could* more progressively price, if only they knew their customers’ WTP more accurately, and yet a more accurate estimate of WTP would require an unethical invasion of privacy; in these cases, firms clearly must abide by such privacy-related moral side constraints and laws.

Second-Order Effects of Unequal Pricing

As was discussed in the introduction, some consumers have strong negative reactions to being treated in ways they perceive as unfair vis-à-vis pricing. These negative reactions are likely associated with negative utility, and yet our Progressive Pricing discussion thus far has not included an endogenous consideration of the likely negative utility some consumers would experience when they learn they are paying more than others. More concretely, in calculating Alice’s utility surplus, we only included her utility gained from the subscription to the *Journal* and her utility lost from the money she spent on it when comparing a Unitary Pricing regime to a Progressive Pricing regime, when it could be argued we should have included some disutility from her perception of being treated unfairly (assuming the *Journal* were transparent about the discrepancy or she found out some other way). In other words, her net utility equation could have been $\text{net utility} = [\text{utility gained from subscription}] - (\text{utility lost from spending } \$X) - (\text{utility lost from unfairness perception})$. This last variable would be zero in the Unitary Pricing regime and (arguably) nonzero in the Progressive regime (at least for some consumers).

This disutility should be factored in, implying that sometimes differences in price levels may need to be attenuated to be less extreme than would otherwise be required were it not for these consumer reactions. In other words, if WTP differs by $10\times$, Progressive Pricing would recommend prices that differ by $10\times$ —however, after factoring in the disutility from differential pricing some would experience, this $10\times$ may need to be reduced to, say, $5\times$.

However, firms are in a position to influence these perceptions. For one, given that Progressive Pricing is superior in the ways we have shown, firms can appeal to this fact in

assuaging consumers, since consumers already widely support price discrimination they view as fair (e.g., senior and student discounts). Firms can also appeal to the fact that, for any individual consumer, while there are likely lower-WTP consumers who are paying less than them, there are also likely *higher*-WTP consumers who are paying *more* than them. Given that all consumers are charged (let’s presume) with the same rule, e.g., 50% of maximum WTP, they are being treated the same in the way that matters. Additionally, we know that as a form of price discrimination becomes more familiar, it is perceived as less unfair. For example, airlines massively price discriminate, Uber has ‘surge pricing’ that was originally decried by riders, hotels charge based on a ‘yield management’ practice, etc.—and yet these practices are widely accepted (if grumbled about), unlike the reactions to Amazon and Staples mentioned above, which consumers do not (yet) find fair. As Progressive Pricing becomes more common, its lay fairness perception will likely increase as well.

Why Would Firms Choose Progressive Pricing?

As we briefly alluded to earlier, our view is that even if firms could estimate an individual consumer’s WTP, they should not then set that individual’s price *equal* to their WTP. Indeed, throughout the paper, we have argued that Progressive Pricing should entail pricing as some *fraction* of an individual consumer’s WTP (e.g., 50%). There are (at least) two reasons why firms should opt to do this.

First, even the most sophisticated algorithms will not be able to estimate an individual’s WTP with complete accuracy. A consumer’s WTP is dependent on a number of constantly varying, context-specific variables (e.g., time of day, psychological disposition, awareness of certain facts in the world, etc.), and so any pricing algorithm will necessarily be tracking a moving target. Furthermore, even if WTP were a static variable, it would still be difficult to measure algorithmically as individuals themselves are often unaware of their own WTP. This means that any estimate of a consumer’s WTP will necessarily have some confidence interval around it that takes into account not only the algorithm’s uncertainty about its own internal estimate, but also the consumer’s uncertainty regarding their own WTP. For example, imagine Firm X’s 95% confidence interval for its estimate of Bob’s maximum WTP for their widget is between 80 and 120. Importantly, there is asymmetric risk to overpricing versus underpricing. Imagine Bob’s actual maximum WTP at the time of sale is 90; if the firm were to set a price of 100 (the middle of their confidence interval), they lose out on the sale completely. Instead, if they had charged at the bottom of their 95% confidence interval, 80, they would have made the sale and only left 10 on the table.

The second reason firms should not seek to price equal to 100% of each consumer's estimated WTP is simply that it is not the right thing to do. That is, when all the consumer surplus is extracted by the producer, the benefits to Progressive Pricing we have argued for across the four consequentialist SWFs are significantly dampened (if not completely eliminated), as total consumer utility surplus goes to 0.

Indeed, in a recent paper, Steinberg (2020) argues that while personalized pricing may lead 'to better overall social welfare...these advantages...are undermined when [it] does not contribute to any socially desired end (other than improved market efficiency in an imperfect market).' He uses a market failures approach, under which 'managers are morally constrained by the point of having competitive markets' (p. 109). He further convincingly argues that when personalized pricing extracts all of the consumer surplus in the market (by charging each consumer their individual maximum WTP), this, 'in a somewhat deontic spirit,' undermines the point of having a market (p. 113). That is, one of the purposes of markets is to compete for consumer surplus, and when firms instead extract all of it for themselves, they deprive the market of a core component of what makes markets beneficial.

However, it is important to underscore that his main argument against price discrimination (although not the only argument he presents) depends upon the premise that '[i]n personalized pricing, each consumer is charged according to their individual reservation price. Personalized pricing extracts all of one's consumer surplus' (p. 98). As we have argued throughout the paper and are emphasizing in this section, however, firms should *not* set each consumer's price equal to their exact maximum WTP. Indeed, his paper helps emphasize the moral (as opposed to just the practical) limitations on price personalization: even if firms could estimate WTP perfectly, charging a price equivalent to 100% of that WTP would be worse than charging as a fraction of that WTP, partly because it eliminates the benefits of Progressive Pricing we espoused previously, but also as it critically undermines a fundamental reason for having markets in the first place, which is to provide consumers with surplus. This 'deontic' reasoning is not meant to override the fundamentally consequentialist framework we have been using, but instead to provide yet another argument against charging 100% of consumers' WTP; that is, it fails both the consequentialist test we have been using with our four SWFs, but also a deontological approach that is encapsulated in the market failures approach Steinberg uses.

Conclusion

Charging customers as a function of their willingness-to-pay (WTP), what we call Progressive Pricing, improves social welfare better than Unitary Pricing. This is true

regardless of which Social Welfare Function (SWF) one holds to be most plausible. Indeed, Progressive Pricing outperforms Unitary Pricing under the 'extreme' SWFs of Utilitarianism (which has no consideration of equality) and Egalitarianism (which has no consideration of total utility), as well as SWFs that will tend to value both more equal and higher total utility distributions (Prioritarianism and Leximin).

This conclusion does not merely imply that price personalization is allowable or morally neutral, as others have argued, but rather that price discrimination is actually *better* for society than Unitary Pricing from a broadly consequentialist standpoint, at least when it is done Progressively. Additionally, Progressive Pricing is the most 'natural' form of price personalization, given its goal is to charge those with higher WTP more than those with lower WTP. This is important, as the classical welfare economics framework has us estimate the impacts on utility and equity separately, with the general assumption they are at odds and must be traded off against each other during evaluation. However, we have shown that Progressive Pricing increases both utility and equity.

Although we have used a relatively innocuous example in this paper (a subscription to a business newspaper), the ethical superiority of Progressive Pricing over Unitary Pricing is increased the more is at stake, for example in healthcare, which has higher stakes both in terms of the price points and the utilities entailed. To illustrate, new orally administered anticancer drugs cost on average over \$100,000 (Dusetzina 2016). Zolgensma (now the most expensive drug in the world), which treats spinal muscular atrophy, is a one-time gene therapy with a cost of \$2.1 million that was recently approved in the US. Given the extremely high price points and the importance of patients receiving these drugs, a Progressive Pricing scheme being used over a Unitary one is all the more important; otherwise, as we have shown, there will be less total utility and less equality.

However, given the extreme consumer reactions and the statements of public intellectuals and academics detailed previously, there will likely be pressure on governments and businesses to limit Progressive Pricing as it becomes more widespread. The purpose of this paper is to show such concerns are unwarranted: higher-WTP customers should pay more than lower-WTP customers, subject to other commonly accepted moral side constraints. However, they should not be charged 100% of their WTP for both the practical and moral reasons we have argued above.

Compliance with Ethical Standards

Conflict of interest The authors declare that they have no conflict of interest.

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

- Ayadi, N., Paraschiv, C., & Rousset, X. (2017). Online dynamic pricing and consumer-perceived ethicality: Synthesis and future research. *Recherche et Applications en Marketing*, 32(3), 49–70.
- Baumol, W. J., & Swanson, D. G. (2003). The new economy and ubiquitous competitive price discrimination: Identifying defensible criteria of market power. *Antitrust Law Journal*, 70(3), 661–685.
- Cohen, G. A. (1989). On the currency of egalitarian justice. *Ethics*, 99, 906–944.
- Dusetzina, S. B. (2016). Drug pricing trends for orally administered anticancer medications reimbursed by commercial health plans, 2000–2014. *JAMA Oncology*, 2(7), 960–961.
- Elegido, J. M. (2011). The ethics of price discrimination. *Business Ethics Quarterly*, 21(4), 633–660.
- Elegido, J. M. (2015). The just price as the price obtainable in an open market. *Journal of Business Ethics*, 130(3), 557–572.
- Huppertz, J. W., Arenson, S. J., & Evans, R. H. (1978). An application of equity theory to buyer-seller exchange situations. *Journal of Marketing Research*, 15(2), 250–260.
- Krugman, P. (2000). What price fairness? *The New York Times* (October 4).
- Levine, M. E. (2002). Price discrimination without market power. *Yale Journal on Regulation*, 19, 1–36.
- Layard, R., Mayraz, G., & Nickell, S. (2008). The marginal utility of income. *Journal of Public Economics*, 92(8–9), 1846–1857.
- Locke, J. (1661/2004). *Venditio*. In *Locke: Political Writings*, David Wooton (ed.), pp. 442–446.
- Martins, M. (1995). An experimental investigation of the effects of perceived price fairness on perceptions of sacrifice and value. *Unpublished doctoral dissertation*, University of Illinois at Urbana-Champaign, Champaign, IL.
- Marcoux, A. M. (2006). Much ado about price discrimination. *Journal of Markets and Morality*, 9(1), 57–69.
- Moriarty, J. (2017). Business ethics. The Stanford Encyclopedia of Philosophy, Edward N. Zalta (ed.). Retrieved from <https://plato.stanford.edu/archives/fall2017/entries/ethics-business/>. Accessed 1 Sept 2019.
- Ordonez, L. D., Connolly, T., & Coughlan, R. (2000). Multiple reference points in satisfaction and fairness assessment. *Journal of Behavioral Decision Making*, 13(3), 329.
- Oswald, A. J. (2005). On the common claim that happiness equations demonstrate diminishing marginal utility of income. IZA Discussion Paper No. 1781. Retrieved from SSRN: <https://ssrn.com/abstract=826388>. Accessed 1 Sept 2019.
- Parfit, D. (1998). Equality and priority. In A. Mason (Ed.), *Ideals of equality* (pp. 1–20). Oxford: Basil Blackwell.
- Phlips, L. (1983). *The economics of price discrimination*. London: Cambridge University Press.
- Rawls, J. (1971). *A theory of justice*. Cambridge, MA: Harvard University Press.
- Schmalensee, R. (1981). Output and welfare implications of monopolistic third-degree price discrimination. *American Economic Review*, 71, 242–247.
- Sen, A. (1980). Equality of what? In S. McMurrin (Ed.), *The Tanner lectures on human values*. London: Cambridge University Press.
- Sachs, J. D. (2017). The efficiency-equity tradeoff. In B. Frey & D. Iselein (Eds.), *Economic ideas you should forget*. New York: Springer.
- Shapiro, C., & Varian, H. R. (1999). *Information rules: A strategic guide to the network economy*. Boston: Harvard Business School Press.
- Snyder, J. (2009a). What's the matter with price gouging? *Business Ethics Quarterly*, 19(2), 275–293.
- Snyder, J. (2009b). Efficiency, equity, and price gouging: A response to Zwolinski. *Business Ethics Quarterly*, 19(2), 303–306.
- Steinberg, E. (2020). Big data and personalized pricing. *Business Ethics Quarterly*, 30(1), 97–117.
- Temkin, L. S. (2000). Equality, priority and the levelling down objection. In M. Clayton & A. Williams (Eds.), *The ideal of equality*. New York: Palgrave MacMillan.
- 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.
- Varian, H. R. (1985). Price discrimination and social welfare. *American Economic Review*, 75, 870–875.
- Zuiderveen Borgesius, F., & Poort, J. (2017). Online price discrimination and EU data privacy law. *Journal of Consumer Policy*, 40, 347.
- Zwolinski, M. (2008). The ethics of price gouging. *Business Ethics Quarterly*, 18(3), 347–378.
- Zwolinski, M., & Wertheimer, A. (2017). Exploitation. The Stanford Encyclopedia of Philosophy, Edward N. Zalta (ed.). Retrieved from <https://plato.stanford.edu/archives/sum2017/entries/exploitation/>. Accessed 1 Sept 2019.

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