

ISSN 1936-5349 (print)
ISSN 1936-5357 (online)

HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS

ALGORITHMIC PRICE DISCRIMINATION
WHEN DEMAND IS A FUNCTION OF BOTH PREFERENCES AND (MIS)PERCEPTIONS

Oren Bar-Gill

Forthcoming in *University of Chicago Law Review*, Vol. 86

Discussion Paper No. 956

05/2018

Harvard Law School
Cambridge, MA 02138

This paper can be downloaded without charge from:

The Harvard John M. Olin Discussion Paper Series:
http://www.law.harvard.edu/programs/olin_center

The Social Science Research Network Electronic Paper Collection:
<https://ssrn.com/abstract=3184533>

Algorithmic Price Discrimination

When Demand is a Function of Both Preferences and (Mis)perceptions

Oren Bar-Gill*

May 2018

Abstract

Sellers are increasingly utilizing big data and sophisticated algorithms to price discriminate among customers. Indeed, we are approaching a world, where each consumer will be charged a personalized price for a personalized product or service. Is this type of price discrimination good or bad? The normative assessment, I argue, depends on the target of discrimination. Sellers are interested in the consumer's willingness-to-pay (WTP) for their goods or services: They maximize profits by charging a price that is as close as possible to the consumer's WTP. This WTP is a function of consumer preferences on the one hand and consumer (mis)perceptions on the other hand. When algorithmic price discrimination targets preferences, it harms consumers but increases efficiency. When price discrimination targets misperceptions, specifically demand-inflating misperceptions, it hurts consumers even more and might also reduce efficiency. In such cases, legal intervention may be needed. In particular, when sellers use personalized pricing, regulators should fight fire with fire and seriously explore the potential of personalized law.

* William J. Friedman and Alicia Townsend Friedman Professor of Law and Economics, Harvard Law School. For helpful comments and suggestions, I would like to thank Omri Ben-Shahar, Yochai Benkler, Ryan Bubb, Glenn Cohen, Hanocho Dagan, Avihay Dorfman, Noah Feldman, Meirav Furth, Assaf Hamdani, Howell Jackson, Louis Kaplow, Alon Klement, Roy Kreitner, Tamar Kricheli-Katz, Adi Leibowitz, Adi Libson, Da Lin, Ariel Porat, Mark Ramseyer, Steve Shavell, Holger Spamann, Cass Sunstein, Doron Teichman, Eyal Zamir, and workshop and conference participants at Harvard, Hebrew University, NYU, Tel-Aviv University, the University of Chicago and the annual meeting of the Israeli Law and Economics Association.

INTRODUCTION

To maximize profits, sellers like to engage in price discrimination – to set higher prices for consumers who are willing to pay more and lower prices for consumers who are willing to pay less. In the past, such price discrimination was limited to coarse categories, e.g., setting higher prices for business travelers and lower prices for leisure travelers. No longer. Fueled by big data, algorithmic price discrimination is able to parse the population of potential customers into finer and finer subcategories – each matched with a different price. In some cases, sellers are even able to set personalized pricing, marching down the demand curve and setting a different price for each consumer.¹

Uber, Amazon, Staples and videogame store Steam were found to vary price by geographic location and, in Uber’s case, also by the time of day. B&Q, a British multinational, tested digital price tags, in its brick-and-mortar stores, that interfaced with customers’ phones and adjusted the displayed price based on the customer’s loyalty cards data and spending habits. Grocery stores are experimenting with digitized and personalized pricing, using e-coupons. Allstate was criticized for optimizing prices based on its calculated likelihood that individual users would comparison-shop before purchasing insurance. A recent study found that many retailers and travel sites set personalized prices that vary by hundreds of dollars from one consumer to the next. Indeed, new players, intermediaries, are appearing with the express purpose of identifying consumers’ willingness to pay and selling this information to retailers who then use it to price discriminate.²

Price discrimination is based on the ability to identify the consumer’s willingness to pay (WTP). Indeed, the concept of WTP is central to any discussion of price discrimination. And yet the richness and subtlety of this concept has not been fully recognized. Standard accounts of price discrimination assume, often implicitly, that a consumer’s WTP is derived entirely from that consumer’s preferences (and budget constraint). For example, some consumers are willing to pay more for a health club membership, because they enjoy a greater benefit from the membership. These consumers may be entertainers whose careers depend on a well-sculpted body or they may enjoy a direct benefit from working out or they enjoy the status that comes with membership in a prestigious health club. Other consumers have different preferences, enjoy a smaller benefit from a health club membership and thus are willing to pay less for the membership. WTP, in the standard model, is all about preferences.³

¹ See Alessandro Acquisti, Curtis Taylor and Liad Wagman, “The Economics of Privacy,” 54 *J. of Econ. Lit.* 442, 466 (2016) (“Tracking and measurability, in addition to websites’ ability to dynamically update and personalize prices for each visitor, are bringing online markets closer to the theoretical scenario of first-degree price discrimination.”) This Article focuses on price discrimination that is based on a consumer’s willingness to pay (WTP). It does not engage with race-based or gender-based discrimination, including price discrimination, that is unrelated to the consumer’s WTP. Note that, in some cases, WTP is correlated with race and gender, and the resulting price discrimination can be analyzed using the model that is developed in this Article.

² See Part I below.

³ See, e.g., Andreu Mas-Colell, Michael D. Whinston and Jerry R. Green, *Microeconomic Theory*, 6–9, 41–63 (OUP, 1995) (explaining that consumers’ individual preferences, described as utility functions, determine

This view of WTP, I argue, is too narrow. In addition to preferences, a consumer's WTP is affected by the consumer's perceptions or misperceptions.⁴ Consider, again, the health club example. Assume that Consumer A and Consumer B hold identical preferences and enjoy an identical benefit – which can be based on health, aesthetics or status – from attending the health club. Further assume that both A and B will attend the club no more than once a month, which inevitably limits the benefit that they gain from membership. Yet, while A realizes that she will attend the club no more than once a month, B falsely believes that he will attend at least once a week and thus overestimates the benefit from membership. This misperception will have a substantial effect on B's WTP; he would be willing to pay more for the health club membership than A would. Even consumers with identical preferences (and budget constraints) may have very different WTP measures. WTP is a function of both preferences and (mis)perceptions. Sellers use big data and sophisticated algorithms to identify and measure both of these factors, so that they can set a personalized price equal to the consumer's full, though misguided, WTP.

A more complete understanding of WTP as a product of both preferences and (mis)perceptions has significant positive and normative implications. In the standard model, where WTP reflects only preferences, price discrimination harms consumers, but increases efficiency. Consumers are harmed, since the seller – a monopolist or any seller with sufficient market power to engage in price discrimination – extracts the entire surplus by setting a price that is just below each consumer's WTP. Efficiency is increased, because price discrimination eliminates the monopoly deadweight loss. The overall welfare assessment depends on the relative weights of efficiency and distribution in the social objective function.⁵ (Unless the adverse distributional implications can be neutralized through an appropriate adjustment to the tax and transfer system.⁶) When WTP reflects both preferences and misperceptions, specifically demand-inflating misperceptions, price

their optimal consumption levels); Hal R. Varian, *Intermediate Microeconomics*, pp. 34–35, 54–55, 73, 78 (8th edition, W.W. Norton and Company, 2010) (same).

⁴ This Article adopts a broad definition of “misperception,” including any mistake that affects the consumer's WTP.

⁵ See, e.g., Mas-Colell et al, *supra* note 3 at 386–87 (“Of course, the distributional properties of this outcome would not be terribly attractive in the absence of wealth redistribution: The monopolist would get all the aggregate surplus generated by its product, and each consumer *i* would receive a surplus of zero.”); Varian, *supra* note 3, at 462–64, 639–640 (describing how perfect price discrimination enables a monopolist to “appropriate [consumer surplus] for itself”); Alessandro Acquisti, Curtis Taylor and Liad Wagman, “The Economics of Privacy,” 54 *J. of Econ. Lit.* 442, 447, 452, 454–457, 466–67 (2016); Curtis R. Taylor, “Consumer Privacy and the Market for Consumer Information,” 35 *RAND J. Econ.* 631, 643 (2004) (finding that personalized pricing can have disparate effects on consumer surplus and total welfare). See also Executive Office of the President, “Big Data: Seizing Opportunities, Preserving Values,” p. 8 (May 2014) (referencing a Council of Economic Advisors study, which found that differential pricing “will often, though not always, be welfare-enhancing” but “can raise concerns about fairness.”); Natasha Singer, *The Government's Consumer Data Watchdog*, *New York Times*, May 23, 2015 (quoting Hal R. Varian, chief economist of Google and professor emeritus at the School of Information at the University of California, Berkeley, about price discrimination: ““It is largely beneficial,” Mr. Varian told the audience, citing examples like discount offers for seniors. He added, “You charge higher prices to people who can afford to pay higher prices.””)

⁶ See Louis Kaplow and Steven Shavell, “Why the Legal System Is Less Efficient than the Income Tax in Redistributing Income,” 23 *J. Legal. Stud.* 667 (1994).

discrimination hurts consumers even more and might also reduce efficiency. The harm to consumers increases, since consumers pay a price that is equal to their perceived benefit, which exceeds their actual benefit. While in the standard model consumers are left with no gain, here consumers suffer an actual loss. Efficiency is also compromised, because price discrimination combined with overestimation leads marginal consumers to purchase products when the cost of production exceeds the actual benefit (but not the higher, perceived benefit). With a larger distributional cost and a smaller efficiency benefit, or even an efficiency cost, the case for legal intervention becomes much stronger.⁷

These results are established using a graph-based, economic analysis. Starting with the neoclassical model, where WTP is derived from preferences alone, I replicate standard results about the effects of price discrimination on the consumer surplus and on overall efficiency. I then turn to the main contribution of this Article: I develop a behavioral model, where WTP is derived from both preferences and misperceptions. In this model, the actual demand curve is supplemented by a perceived demand curve. Specifically, I study the implications of an overestimation bias that pushes the perceived demand curve above the actual demand curve. (I argue that a demand-inflating, overestimation bias is more likely in this context; the implications of underestimation are explored in the appendix.) I show that the overestimation bias can substantially increase the welfare costs of algorithmic price discrimination. (The analysis compares no price discrimination to perfect price discrimination. I omit the qualifier “perfect” for the sake of brevity. Most of the results hold also with imperfect price discrimination.)

While this Article focuses on the costs of big data and the algorithms that utilize these data, it is important to emphasize that algorithms wielding big data also produce substantial benefits. They allow for more precise matching between consumers and the products and services that these consumers need or want. The online ads that are relevant for one consumer would be annoying spam for another. Of course, matching is only efficient, when the consumer is offered a product that she really need or wants (preferences), not when she is offered a product that she mistakenly thinks she needs or wants (misperceptions).⁸ Big

⁷ See also Paul Heidhues and Botond Koszegi, “Naivete-Based Discrimination,” *Quar. J. Econ.* (2016) (arguing that price discrimination that targets consumers naivete, a type of misperception, reduces welfare, albeit for reasons that are different than those analyzed here); ARIEL EZRACHI, ARIEL AND MAURICE E. STUCKE, *VIRTUAL COMPETITION: THE PROMISE AND PERILS OF THE ALGORITHM-DRIVEN ECONOMY*, p. 117 (2016) (arguing that “behavioural discrimination” – the authors’ term for discrimination that targets cognitive biases – represents a starker wealth transfer from consumers to producers; specifically, “[t]his behavioral discrimination can increase consumption, optimize the extraction of wealth, and affect other important values, such as privacy, equality and fairness.”) *See generally id.* at 117–130.

⁸ Another example of undesirable matching occurs when sellers vary the quality of an offered product or service based not on the consumer’s preferences, but on other factors that make the consumer less profitable to the seller or reduce the consumer’s threat to switch to a competitor. For example, CableOne bragged that it offers worse service to customers with low FICO scores. See: Karl Bode, *CableOne Brags It Provides Worse Service to Bad Credit Customers*, *DSL Reports*, May 27, 2016 (<http://www.dslreports.com/shownews/CableOne-Brags-It-Provides-Worse-Service-to-Bad-Credit-Customers-137070>). Cf. David Gilo and Ariel Porat, “The Hidden Roles of Boilerplate and Standard-Form Contracts: Strategic Imposition of Transaction Costs, Segmentation of Consumers, and Anticompetitive Effects,” 104 *Mich. L. Rev.* 983 (2006) (discussing other costs of quality-based discrimination, focusing on contractual terms as a quality dimension).

data and sophisticated algorithms can also prevent inefficient cross-subsidization. For example, in the consumer credit market the seller's cost depends on the individual consumer's creditworthiness. In such markets, big data enables a different type of personalized pricing – pricing that tracks the seller's cost, rather than the consumer's WTP. This form of price discrimination is much less objectionable and, in fact, is likely to increase efficiency and help consumers.⁹ Finally, it is important to remember that big data and algorithms also bring us price comparison sites and apps, which work against price discrimination.

Since price discrimination is more harmful with misperception, we should consider policies that curb price discrimination (even if we were uncertain about such policies in a world without misperception). Policymakers can target the foundation of algorithmic price discrimination: big data. Since algorithmic price discrimination is fueled by big data, this extreme form of price discrimination can be curbed by limiting sellers' access to information about consumers' WTP. The analysis thus provides a new argument for increased privacy protections and enhanced data security measures. But attacking the big data foundation of price discrimination runs the risk of throwing the baby out with the bathwater. Given the benefits that personalization provides, cutting the flow of information might be a net loss for consumers.

Personalized law suggests a more promising set of responses. Demand-inflating misperceptions result in high prices. A standard response is to impose price caps.¹⁰ But standard price caps are homogeneous, a single cap for the entire market, whereas price discrimination results in prices that are too high but also heterogeneous. If a single cap doesn't work, how about multiple, personalized price-caps? If sellers use personalized pricing, then regulators can respond with personalized price caps.¹¹ The personalized cap could track the misperception and thus undo its effects. Another, softer approach uses the common disclosure technique, but with a new twist. In a world of personalized pricing – the world of price discrimination – disclosure mandates should also be personalized. Namely, sellers should be made to provide each consumer with individualized information on the benefit that the particular consumer stands to gain from the product. With such information, consumers will be able to separate out the misperception-based component of their WTP and avoid overpaying for the product or service.

A final comment, on the domain of analysis: While this Article focuses on price discrimination in consumer markets, a similar phenomenon occurs in labor markets. Employers are increasingly using big data and sophisticated algorithms to set personalized

⁹ See, e.g., PYMNTS.com, “What Does Phone Usage Say About Creditworthiness?” Dec. 1, 2015 (<https://perma.cc/U72Q-MGZC>) (several Silicon Valley start-ups have used the large amount of data collected on consumers' phones to assess their creditworthiness and offer them small loans at lower interest rates). See also Wonga.com (online payday lending based on algorithms that utilize Facebook data to predict probability of default); Ariel Ezrachi, Ariel and Maurice E. Stucke 2016. *Virtual Competition: The Promise and Perils of the Algorithm-Driven Economy*, p. 90 (2016) (citing a *Wall Street Journal* report, from 2010, that Capital One personalized its credit card offers based on new users' credit scores).

¹⁰ See, e.g., Oren Bar-Gill, *Price Caps in Multiprice Markets*, 44 *J. Legal Stud.* 453 (2015).

¹¹ The idea to use personalized price caps was first suggested by Ziv Ben-Shahar. See Ziv Ben-Shahar, *Ethics and Regulation of Personalized Pricing*, Working Paper (2017).

wages just above the individual employee's Willingness to Accept (WTA).¹² Here too the problem can become more severe when we look beyond the standard, preference-based model and recognize that WTA is a product of both preferences and (mis)perceptions.

The remainder of this Article is organized as follows. Part I sets the stage. Section I.A documents the rise of personalized pricing, fueled by the big data revolution. Section I.B turns to WTP as the main target of algorithmic price discrimination. It demonstrates that misperception constitutes an important component of WTP and offers evidence on the prevalence of misperception in consumer markets and, specifically, on misperception as the target of price discrimination. Part II is the analytical heart of the Article. It derives the positive and normative implications of price discrimination, when WTP is the product of both preferences and misperceptions. Part III discusses policy implications. The Article ends with a brief conclusion. An appendix considers two extensions: the first introduces a positive or negative correlation between preferences and misperceptions, and the second studies demand-deflating misperceptions.

I. PERSONALIZED PRICING, WTP AND CONSUMER MISPERCEPTIONS

A. Big Data and the Rise of Personalized Pricing

Sellers can make more money when they set higher prices for consumers who are able and willing to pay more. Big data and sophisticated algorithms allow sellers to do just that – to set personalized prices. Many sellers collect, or purchase, reams of data about existing and potential customers.¹³ And a growing number of sellers, across multiple markets, are

¹² Employers have long screened applications based on credit reports and background checks, but many now also research a candidate's presence online. See Jennifer Preston, *Social Media History Becomes a New Job Hurdle*, N.Y. TIMES, July 21, 2011, at B1; Cross-Tab, *Online Reputation in a Connected World* 6 (Jan. 2010), <http://go.microsoft.com/?linkid=9709510> (reporting that 75% of U.S. employers research applicants online). See generally Ian Byrnsie, Note, *Six Clicks of Separation: The Legal Ramifications of Employers Using Social Networking Sites to Research Applicants*, 10 VAND. J. ENT. & TECH. L. 445 (2008); Nathan J. Ebnet, Note, *It Can Do More than Protect Your Credit Score: Regulating Social Media Pre-Employment Screening with the Fair Credit Reporting Act*, 97 MINN. L. REV. 306 (2012).

¹³ See, e.g., Liz Mineo, *On Internet Privacy, Be Very Afraid*, Harvard Gazette, August 25, 2017 (quoting cybersecurity expert, Bruce Schneier: "I used to say that Google knows more about me than my wife does, but that doesn't go far enough."); Sandy Parakilas, *We Can't Trust Facebook to Regulate Itself*, N.Y. Times, Nov. 19, 2017 (Sandy Parakilas, a former operations manager at Facebook, described its former employer as "a company that reaches most of the country every day and has the most detailed set of personal data ever assembled."); Federal Trade Commission, *Data Brokers: A Call for Transparency and Accountability* 23–31 (May 2014) (describing how data brokers sell retailers a variety of marketing products, with the top five brokers alone generating revenue of nearly \$200 million in 2012); Akiva A. Miller, *What Do We Worry about When We Worry about Price Discrimination - The Law and Ethics of Using Personal Information for Pricing*, 19 J. TECH. L. & POL'Y 41, 45–47 (2014) (listing a number of techniques that online and brick-and-mortar retailers use to capture customer data).

making use of data analytics to identify consumers' WTP and to adjust pricing based on the individual consumer's WTP.¹⁴

Sellers price discriminate based on several types of information, including consumers' location, the time of day, the characteristics of the consumer's computer (e.g., operating system and browser) and their purchase history.¹⁵ Sellers can collect this information themselves or purchase it from data brokers. Hannak et al. analyzed real-world data from 300 consumers who visited 16 e-commerce websites, and found "evidence of personalization on four general retailers and five travel sites, including cases where sites altered prices by hundreds of dollars."¹⁶

Evidence suggests that a consumer's geographic location affects sellers' pricing decisions. For example, Uber collects rich geographic data about its customers and uses it to price discriminate.¹⁷ Also, Amazon, Staples and videogame store Steam were found to vary price by geographic location by as much as 166%.¹⁸ Uber, reportedly, uses time, in addition to location, to price discriminate, calculating a rider's WTP for a particular route at a certain time of day and adjusting prices accordingly.¹⁹ Characteristics of the consumer's computer, tablet or smartphone have also been shown to trigger differential pricing. Generally, Apple iOS and Safari users pay higher prices.²⁰ Sophisticated algorithms have identified correlations between these (and other) factors and consumers' WTP and use these correlations to guide seller's pricing strategies.

Perhaps most useful for the price-discriminating seller is a consumer's purchase history. Purchase history, and even browsing history, on the seller's website or on other e-commerce sites, is an obvious source of WTP information for pricing algorithms. B&Q, a

¹⁴ See Executive Office of the President, "Big Data: Seizing Opportunities, Preserving Values," p. 8 (May 2014) (noting that companies increasingly use big data for personalized pricing); Akiva A. Miller, What Do We Worry about When We Worry about Price Discrimination - The Law and Ethics of Using Personal Information for Pricing, 19 J. TECH. L. & POL'Y 41 (2014) (sellers are increasingly able "to identify users and collect personal data, target advertising, and adjust prices rapidly and automatically."); Adam Tanner, Different Customers, Different Prices, Thanks to Big Data (available at <http://www.forbes.com/sites/adamtanner/2014/03/26/different-customers-different-prices-thanks-to-big-data/#5f621414f31c>) (<https://perma.cc/Q6WY-V25U>); Zuiderveen Borgesius, F.J. & Poort, J., "Online Price Discrimination and EU Data Privacy Law," J. Consum. Policy (2017).

¹⁵ See Miller, *supra* note 5; Aniko Hannak et al., Measuring Price Discrimination and Steering on E-commerce Web Sites, <http://www.ccs.neu.edu/home/cbw/static/pdf/imc151-hannak.pdf> (last visited Feb 26, 2017).

¹⁶ See: Aniko Hannak et al., Measuring Price Discrimination and Steering on E-commerce Web Sites, <http://www.ccs.neu.edu/home/cbw/static/pdf/imc151-hannak.pdf> (last visited Feb 26, 2017)

¹⁷ See Kominers (2017) on price discrimination by Uber (<https://perma.cc/2T8E-J5K8>).

¹⁸ See: Jakub Mikians, László Gyarmati, Vijay Erramilli, Nikolaos Laoutaris, Detecting Price and Search Discrimination on the Internet, Proceedings of the 11th ACM Workshop on Hot Topics in Networks, 2012.

¹⁹ Eric Newcomer, *Uber Starts Charging What It Thinks You're Willing to Pay*, Bloomberg Technology, 19 May 2017, <https://bloom.bg/2qBCuyp> ("Uber calculates riders' propensity for paying a higher price for a particular route at a certain time of day. For instance, someone traveling from a wealthy neighborhood to another tony spot might be asked to pay more than another person heading to a poorer part of town, even if demand, traffic and distance are the same.")

²⁰ See: Aniko Hannak et al., Measuring Price Discrimination and Steering on E-commerce Web Sites, <http://www.ccs.neu.edu/home/cbw/static/pdf/imc151-hannak.pdf> (last visited Feb 26, 2017)

British multinational, tested digital price tags, in its brick-and-mortar stores, that interfaced with customers' phones and adjusted the displayed price based on the customer's loyalty cards data and spending habits.²¹ Grocery stores are experimenting with personalized pricing, using digital coupons.²² And Allstate was criticized for optimizing prices based on its calculated likelihood that individual users would comparison-shop before purchasing insurance.²³ Such examples of personalized pricing are becoming increasingly prevalent.

Indeed, new players, intermediaries, are appearing with the express purpose of identifying consumers' WTP and selling this valuable information to retailers. For example, a start-up company, Optimal Decisions Group, conducted research on consumers' WTP and sold this research, together with profit-maximizing pricing models, to insurers. Another start-up, Freshplum, developed algorithms that are able to predict which customers will buy only if given a discount and which are willing to pay higher prices, and sold this targeted-discounts information to e-commerce vendors. And Tellapart, operating in the music industry, collects consumer data and advises sellers on how to personalize their pricing schemes.²⁴

While algorithmic price discrimination is clearly on the rise, it is important to recognize several important limits to such personalized pricing. The first is arbitrage. When a consumer with a low WTP can turn around and resell the product to a consumer with a high WTP, then price discrimination becomes more difficult.²⁵ Personalized pricing is, therefore, more likely to occur in markets, where arbitrage opportunities are limited, e.g., services markets, markets for perishable goods and markets where high transaction costs (including search costs) make arbitrage difficult.²⁶ A second limit on personalized pricing is fairness. A consumer who learns that she paid much more than another consumer for the

²¹ Sal Thomas, *Does dynamic pricing risk turning personalisation into discrimination?*, October 22, 2014, <http://bit.ly/2yCAvxP>.

²² PYMNTS, "In-Store Tracking Tech Gets Personalized," Feb. 9, 2018 (noting the trend towards digitized and personalized pricing in the grocery market, using coupons).

²³ See Ezrachi and Stucke, *Virtual Competition*, supra note 7, at 90.

²⁴ See Tanner, supra note 5. See also Kyle James, "Abandon Shopping Cart: 17 Online Retailers Who'll Bait You Back with a Coupon," rather-be-shopping.com, July 18, 2014 (identifying 17 retailers that offered coupons to customers who abandoned their carts, including Toys R' Us, Bass Pro Shops, Bed Bath & Beyond, Dick's Sporting Goods, JC Penney, Macy's, OfficeMax, and Williams-Sonoma). The music industry example, Tellapart, is discussed in: Adi Libson and Gideon Parchomovsky, *Toward the Personalization of Copyright* [in this Symposium] (referring to TELLAPART, *Predictive Marketing Just Got Personal*, <https://www.tellapart.com>; and, with respect to a business Tellapart took over – Freshplum – and other competitors, to: Adam Tanner, *Different Customers, Different Prices, Thanks to Big Data*, FORBES, April 14, 2014, available at <https://www.forbes.com/sites/adamtanner/2014/03/26/different-customers-different-prices-thanks-to-big-data/#5dc657545730>.)

²⁵ This assumes that sellers have information only on the consumer's use-based WTP, namely, the WTP that applies if the consumer purchases the product for personal use. For a consumer who plans to resell the product, the true WTP is higher than the use-based WTP. If sellers have information on the true WTP, then the arbitrage limit disappears.

²⁶ See, e.g., Lars. A Stole, "Price Discrimination and Competition," in Mark Armstrong and Robert H. Porter (eds.), *Handbook of Industrial Organization*, Vol. 3, Ch. 34, p. 2226 ("It is well known that price discrimination is only feasible under certain conditions: (i) firms have short-run market power, (ii) consumers can be segmented either directly or indirectly, and (iii) arbitrage across differently priced goods is infeasible.")

exact same product would feel wronged; and such outrage is bad for business.²⁷ Personalized pricing is more likely to occur in markets where consumers do not have direct information about how much other consumers are paying and in markets where differential pricing is coupled with heterogeneous products such that comparisons between consumers become difficult.²⁸

B. WTP: Between Preferences and (Mis)Perceptions

We have seen that sellers increasingly collect and utilize big data and sophisticated algorithms to identify consumers' WTP and set personalized prices equal or close to the individual consumer's WTP. But what exactly is WTP? What is it comprised of? What determines the maximum price that a consumer is willing to pay for a product or service? In the standard, neoclassical model, WTP is derived from the consumer's preferences. Preferences are heterogeneous. Some consumers gain a large benefit from a certain product and are thus willing to pay more for that product. Other consumers gain a smaller benefit and are willing to pay less for the same product. Consider a high-end steak dinner. Some consumers really love steak and are willing to pay a lot for the dinner. Other consumers enjoy steak, but equally enjoy a much cheaper burger. These consumers would be willing to pay less for the steak dinner. And a third group of vegetarian consumers have a WTP of zero. (The consumer's budget constraint also plays an important role in determining her WTP. But since the budget constraint plays an equally important role in the behavioral model described below, it is not emphasized here.²⁹)

The standard, preference-based theory of WTP is clearly right and clearly incomplete. Preferences surely affect WTP, but other forces have an equally important, and sometimes

²⁷ See Daniel H. Kahneman, Jack L. Knetsch and Richard H. Thaler, "Fairness As a Constraint on Profit Seeking: Entitlements in the Market," 76 *Amer. Econ. Rev.* 728 (1986); Eric T. Anderson and Duncan I. Simester, "Price Stickiness and Customer Antagonism," 125 *Quar. J. Econ.* 125 729 (2010); Kominers, *supra* note 7 (successful price discrimination can result in bad publicity and shift business to competitors); Tanner, *supra* note 5 (customers have reacted negatively to stories about differential pricing, leading companies to guard such policies or avoid surcharges in favor of personalized discounts, which consumers accept more); Zuiderveen Borgesius, F.J. & Poort, J., "Online Price Discrimination and EU Data Privacy Law," *J. Consum. Policy* (2017). The importance of the fairness limit is debatable. For example, airlines have long charged substantially different prices for identical seats, without triggering a significant consumer uproar. Either consumers are less sensitive to fairness concerns in this context, or the airlines have managed to keep the differential pricing less salient to most consumers.

²⁸ Consumers' "aversion to surveillance" could impose another, market-based limit on personalized pricing, if the personalized pricing suggests to consumers that they are being surveilled by the seller.

²⁹ Budget constraints and wealth effects are relevant to the analysis in this Article in two, related ways: First, budget constraints and wealth effects suggest a positive correlation between WTP and wealth and are, therefore, relevant for identifying the distributional implications of price discrimination. Second, budget constraints and wealth effects suggest that a consumer's WTP may be lower than the preference-based benefit that the consumer would gain from the product or service. This means that, in the standard model, perfect price discrimination does not leave consumers with zero surplus. And, in the behavioral model, it means that perfect price discrimination may reduce the consumer's gain from the transaction, rather than inflicting an actual loss. At the extreme, if budget constraints rather than preferences determine consumers' WTP, then misperceptions about the benefit from a product or service will have no effect, although misperceptions about price (or about the budget constraint itself) would continue to affect the analysis.

more important, role in determining a consumer's WTP. In particular, a behavioral model of WTP recognizes that perceptions, and misperceptions, substantially affect a consumer's WTP. Reconsider our steak dinner. In addition to preferences for steak, consumers' WTP can be affected by their perceptions about the benefits and risks from eating steak. For example, is the steak organic or not? Does it come from grass-fed or grain-fed cows? Is it free from antibiotics and hormones? How were the cows treated before they were slaughtered? At least some consumers would be interested in the answers to these questions, and their beliefs about the answers will affect their WTP. Moreover, for some (many?) consumers these beliefs will be false. They might think that their steak comes from a grass-fed cow, when in fact it doesn't. Or they might think that "organic" means one thing, when in fact it means something else.³⁰ Perceptions, and misperceptions, clearly affect WTP.

WTP is also affected by perceptions, and misperceptions, about the price charged. Especially when price is contingent and multidimensional, consumers might underestimate the total price that they will end up paying for the product or service. Consider a credit card with an annual fee of \$50 and a late fee of \$25. The consumer's true WTP, based on the benefit that this consumer obtains from using the credit card, is, say, \$60. Assume that this consumer will be late twice, and thus pay \$50 in late fees. The total price for this consumer is \$100 (= \$50 annual fee plus \$50 in late fees). And so, with accurate perceptions, the consumer will not get the credit card. But what if the consumer optimistically thinks that she will never be late. This misperception transforms the total price of \$100 into a perceived total price of \$50, and the consumer will get the card. The misperception operates on the price, and not directly on the consumer's WTP. But the effect is the same. The misperception increases the de facto WTP by \$50.

Misperceptions that affect consumers' WTP are quite common. They have been documented across multiple consumer markets. Consumers overestimate the benefits from health club subscriptions, because they overestimate how often they will attend the health club.³¹ Consumers overestimate the benefits, or underestimate the risks, from different types of foods, because they do not understand the meaning of "organic," "natural," "fruit juice drink" and other enticing labels.³² And consumers underestimate the price that they will pay for credit cards, cell phone plans, mortgages and many other products and services.³³

In most cases, the misperception manifests as an overestimation of value, or underestimation of price. The reason is that sellers and service providers have strong incentives to promote overestimation of value and combat underestimation of value; and, similarly, to promote underestimation of price and combat overestimation of price. Indeed,

³⁰ See Oren Bar-Gill and Kevin Davis, "(Mis)perceptions of Law in Consumer Markets," *Amer. L. and Econ. Rev.* (2017) (citing evidence of consumer misperceptions about the benefits and risks associated with different food products and how such misperceptions affect purchasing decisions).

³¹ See Stefano DellaVigna and Ulrike Malmendier, "Paying Not to Go to the Gym," 96 *Amer. Econ. Rev.* 694 (2006).

³² See Bar-Gill and Davis, *supra*; 21 CFR 102.33(a) (Regulating labels of beverages with fruit content).

³³ See OREN BAR-GILL, *SEDUCTION BY CONTRACT: LAW, ECONOMICS AND PSYCHOLOGY IN CONSUMER MARKETS* (2012).

sellers and service providers design their products, contracts and pricing schemes to create or enhance misperceptions that increase consumers' WTP.³⁴ Therefore, the analysis, in the next section, focuses on the implications of algorithmic price discrimination, when the perceived WTP exceeds the true WTP.³⁵ (In some markets, misperception manifests as an underestimation of value, or overestimation of price, such that the perceived WTP is smaller than the actual WTP. This possibility is explored in the appendix.)

Before turning to that analysis, however, it is important to emphasize the link between price discrimination and consumer misperception. We have seen that sellers increasingly use big data and sophisticated algorithms to identify a consumer's WTP and set personalized prices, namely, to price discriminate. We have also seen that WTP is generally the product of both preferences and (mis)perceptions. It follows that price discrimination will target both preferences and misperceptions as equally important determinants of a consumer's willingness-to-pay.³⁶

Indeed, sellers are striving to improve their understanding of consumer perceptions and misperceptions. Industry publications suggest that big data and sophisticated algorithms are being used to redefine the relationship with the consumer – to create “a different engagement effect with a particular consumer based on what frame of mind they're in at that point in time.”³⁷ The reference to a consumer's “frame of mind” suggests that sellers and service providers are reaching beyond preferences and entering the realm of perceptions and misperceptions. Related accounts suggest that Facebook and Google are using the vast amounts of data at their disposal to implement “cognitive” services, including “assessing someone's personality by sifting through their writings.”³⁸ These services can then be utilized directly by Facebook and Google or offered to other sellers and service providers. Again, the reference to “cognitive” services that assess a consumer's personality suggests that sellers are trying to identify the misperception component of the WTP.

³⁴ Id. On the other hand, there are market forces that work to reduce misperception. For example, if misperception benefits one seller, her competitor may try to educate consumers and reduce the misperception.

³⁵ There are important limits on the prevalence of demand-inflating misperceptions. First, if a seller promotes a misperception that boosts demand for her own product, this misperception may well reduce demand for her competitor's product. Second, since consumers operate within a budget constraint, if all sellers attempt to inflate demand for their respective products, then these attempts will inevitably cancel out on average.

³⁶ See: Paul Heidhues and Botond Koszegi, “Naivete-Based Discrimination,” *Quar. J. Econ.* (2016). Heidhues and Koszegi develop a model where sellers price discriminate based on whether a consumer is sophisticated or naïve (or, more generally, on the consumer's degree of naivete). Naivete can be understood as a type of misperception – a misperception of one's time preferences. One of the leading examples discussed by Heidhues and Koszegi, concerns misperceptions about the likelihood of using, perhaps inadvertently, the overdraft feature of a bank account and thus incurring overdraft fees. The authors explain how banks will offer different prices to different consumers, raising the overdraft fee charged to naïve consumers. See also: Alexei Alexandrov, *Generalizing Lessons from Behavioral Economics Across Many Biases*, Working Paper, Sec. 3 (2015) (available at SSRN: <http://ssrn.com/abstract=2576004>) (studying a general model of perfect price discrimination with a single biased consumer).

³⁷ Actionable Insights: Issuers' 2015 Resolution, PYMNTS, Dec. 30, 2014.

³⁸ See: Fuel of the future: Data is giving rise to a new economy, *The Economist*, May 6, 2017.

As observed by Ryan Calo, “[e]merging methods of big data present a new and vastly more efficient way to identify cognitive bias by attempting to pinpoint profitable anomalies.”³⁹ Calo further suggests that emerging technologies will “empower corporations to discover and exploit the limits of each individual consumer’s ability to pursue his or her own self-interest. Firms will increasingly be able to trigger irrationality or vulnerability in consumers.”⁴⁰ While sellers are evidently interested in consumer psychology, the analysis below does not depend on sellers’ ability to disentangle the preference-based and misperception-based components of the WTP. Indeed, to price discriminate, sellers only need to know the full, albeit misguided, WTP.

II. MARKET OUTCOMES AND WELFARE EFFECTS

My goal is to study the effects of both price discrimination and misperception and, specifically, the interaction between price discrimination and misperception. I, therefore, construct, and study, the following 2x2 matrix:

		Misperception	
		No	Yes
Price Discrimination	No	No Misperception, No Price Discrimination (Figure 1b)	Misperception, No Price Discrimination (Figure 2b)
	Yes	No Misperception, With Price Discrimination (Figure 1c)	Misperception, With Price Discrimination (Figure 2c)

Table 1: Organization of the Analysis

The “No Misperception” column replicates the conventional microeconomic analysis of first-degree price discrimination and how it compares to a no-price-discrimination benchmark.⁴¹ I review this analysis briefly in Section A. Then, in Section B, I turn to the “Yes Misperception” column, where the main contribution of this Article lies. (As

³⁹ Ryan Calo, Digital Market Manipulation, 82 Geo. Wash. L. Rev. 995, 1010 (2014).

⁴⁰ Id. at 999. See also: Ezrachi, Ariel and Stucke, Maurice E., The Rise of Behavioural Discrimination (August 25, 2016). Available at SSRN: <http://ssrn.com/abstract=2830206> (describing ‘behavioural discrimination’ – the combination of price discrimination techniques and behavioral economics to exploit consumers’ cognitive biases and to maximize firm profits by inducing consumers to make irrational purchasing decisions).

⁴¹ See, e.g., Mas-Colell et al, supra note 3 at 384–87; Varian, supra note 3, at 439–449, 462–64. See also Arthur Pigou, THE ECONOMICS OF WELFARE (4th ed., 1932) 161 (coining the term “first-degree price discrimination”).

explained above, I focus initially on demand-inflating misperceptions. Demand-deflating misperceptions are discussed in the appendix.)

A. No Misperception

Consider the most basic market setup with a linear, downward sloping demand curve and a linear, horizontal supply curve (reflecting a fixed-per-unit-cost assumption; let k denote the per-unit cost).⁴² See Figure 1a. In a (perfectly) competitive market, the intersection of the demand curve with the supply curve, at (Q_C, P_C) , represents the market equilibrium, where Q_C represents the equilibrium quantity and P_C represents the equilibrium price (which is equal to the per-unit cost, k).⁴³ In terms of welfare, the red triangle represents the consumer surplus – the difference between the consumer’s WTP and the price, P_C , aggregated across all consumers. (Some consumers have a high WTP. They are represented by the high points on the left side of the demand curve, and they enjoy more surplus. Other consumers have a lower WTP. They are represented by the lower points of the demand curve, close to Q_C , and they enjoy less surplus.) In a competitive market, the seller enjoys no surplus at all – selling her product at a price, P_C , that precisely equals her per-unit cost of production. Therefore, social welfare is equal to the consumer surplus, as represented by the red triangle.

⁴² See, e.g., Mas-Colell et al, *supra* note 3 at 321; Varian, *supra* note 3, at 292-94.

⁴³ See, e.g., Mas-Colell et al, *supra* note 3 at 316–322; Varian, *supra* note 3, at 292-94.




Figure 1a: No Misperception, No Price Discrimination, Perfect Competition

The perfect competition scenario serves as a useful baseline, but a monopolistic market is more relevant for our purposes, since at least some degree of market power is necessary for price discrimination.⁴⁴ (I focus on the extreme, monopoly case for ease of exposition.) To assess the effects of price discrimination, we must first consider, as a reference point, a monopolistic market in the absence of price discrimination, namely when the monopolist sets a single price. As compared to the competition case, a monopolist will set a higher price, $P_M > P_C$, and sell fewer units of the product, $Q_M < Q_C$. See Figure 1b. Consumer surplus is now represented by the smaller red triangle. In a monopolistic market, the seller enjoys a positive surplus, which is represented by the blue rectangle and is equal to the number of units sold multiplied by the difference between the monopoly price and the per-unit cost: $Q_M \cdot (P_M - k)$. Social welfare is, by definition, equal to the sum of the consumer surplus and the producer's (monopolist's) surplus. The black triangle represents the monopoly deadweight loss: Because of the higher price that the monopolist charges, consumers who should buy the product refrain from purchasing it (specifically, the lost quantity is given by $Q_C - Q_M$); and the welfare that these lost purchases would have produced constitutes the monopoly deadweight loss.

⁴⁴ See, e.g., Lars. A Stole, "Price Discrimination and Competition," in Mark Armstrong and Robert H. Porter (eds.), *Handbook of Industrial Organization*, Vol. 3, Ch. 34, p. 2226 ("It is well known that price discrimination is only feasible under certain conditions: (i) firms have short-run market power, (ii) consumers can be segmented either directly or indirectly, and (iii) arbitrage across differently priced goods is infeasible.")




Figure 1b: No Misperception, No Price Discrimination, Monopoly

It is time to introduce price discrimination. Specifically, I now allow the monopolist to charge each consumer a different, personalized price. Using big data and sophisticated algorithms, the monopolist will identify each consumer's WTP and set a personalized price just below this WTP. Thus a consumer with a high WTP on the left side of the demand curve will pay a high price; and a consumer with a lower WTP towards the middle or right side of the demand curve will pay a lower price. See Figure 1c. Observe that the quantity sold is Q_C , as in the competition case. But while competition gives the entire surplus to consumers (the red triangle in Figure 1a), a price discriminating monopolist keeps the entire surplus to itself (the blue triangle in Figure 1c; note that this blue triangle is equal in size to the red triangle in Figure 1a). Price discrimination allows the monopolist to increase the quantity sold – from Q_M to Q_C – thus eliminating the deadweight loss and increasing the overall social welfare. This efficiency gain comes at a steep distributional price, however, as the entire surplus goes to the monopolist and consumers are left with nothing.⁴⁵ Still, the efficiency gain is worth emphasizing. It is a powerful argument in favor of price

⁴⁵ There are additional distributional implications for the consumer side of the market. Consumers with a higher WTP, who would have purchased the product at the (no discrimination) monopoly price, suffer an affirmative loss, as they pay more for the same product. Consumers with a lower WTP are not affected – without price discrimination they would have been priced out of the market, and with price discrimination they still get a zero (net) surplus. Budget constraints and wealth effects (see supra note 29) add nuance to this analysis: If WTP is smaller than the preference-based benefit, then perfect price discrimination while hurting consumers with a high WTP, affirmatively helps consumers with a low WTP. And if WTP is positively correlated with wealth, then perfect price discrimination results in progressive redistribution.

discrimination in the standard, no misperception model. We next turn to the misperception case and show that this efficiency gain shrinks and might even become an efficiency loss.

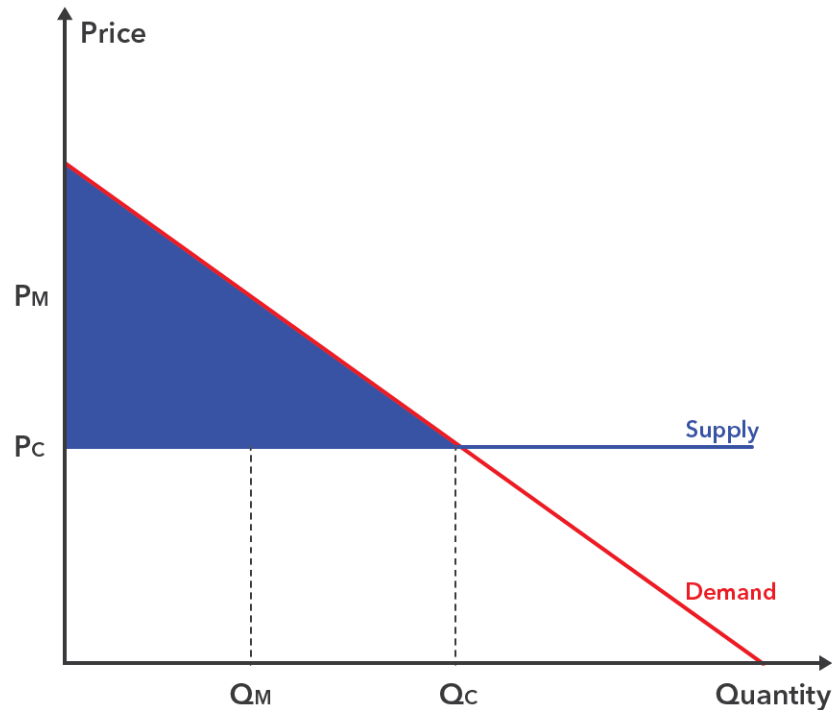


Figure 1c: No Misperception, With Price Discrimination, Monopoly

B. Misperception

Consumers might either overestimate or underestimate the benefit from a product. I focus on overestimation for the reasons noted above – sellers have an incentive to promote overestimation and fight underestimation. I initially assume that the degree of overestimation is not correlated with consumers’ preference-based WTP, namely, that the average bias level is the same for consumers with a higher preference-based WTP at the left-hand side of the demand curve and for consumers with a lower preference-based WTP towards the middle and right-hand side of the demand curve. (This assumption is relaxed in Section C below.) Now, in addition to the actual demand curve, we have a perceived demand curve. In Figures 2a, 2b and 2c, the actual demand curve is represented by the solid red line, and the perceived demand curve is represented by the dashed red line.

I start, as before, with the perfect competition baseline. The price remains $P_C = k$, but the equilibrium quantity, Q_C' , is larger. This quantity is determined by the intersection of the supply curve with the perceived demand curve. See Figure 2a. In welfare terms, consumers in the $[0, Q_C]$ range still enjoy the same consumer surplus, as in the no misperception case. This surplus is represented by the red triangle. These consumers mistakenly believe that they are getting an even larger surplus – the perceived extra surplus is represented by the

light-red trapezoid (above the supply curve) – but this false belief does not harm them. There is, however, another group of consumers who are harmed. These consumers, in the $[Q_C, Q_{C'}]$ range, should not purchase the product. They buy only because of the misperception, because they overestimate the product's value. These purchases create a welfare loss, which is both an efficiency loss and a loss to consumers (namely, the consumers bear the entire efficiency loss). This loss is represented by the light-red triangle (below the supply curve).




Figure 2a: Overestimation, No Price Discrimination, Perfect Competition

Next, consider a monopolistic market, but without price discrimination. In the standard model, without misperception, the monopoly price is determined by the demand curve.⁴⁶ And, when consumers overestimate the value of the product, the price is determined by the perceived demand curve. Therefore, the monopoly price with misperception, P_M' , is higher than the monopoly price without misperception, P_M . The quantity sold with misperception, Q_M' , is also higher than the quantity sold without misperception, Q_M .⁴⁷ See Figure 2b. Turning to welfare: The higher price reduces the actual consumer surplus, which is represented by the red triangle. The *perceived* surplus is larger – the perceived extra surplus is represented by the light-red trapezoid. Overestimation causes some consumers to purchase the product even though its actual value to them is lower than the price, P_M' . The loss incurred by these consumers is represented by the purple triangle. This loss reduces

⁴⁶ See, e.g., Mas-Colell et al, *supra* note 3 at 384–86; Varian, *supra* note 3, at 441-43.

⁴⁷ The overestimation inflates demand and thus increases the quantity sold. The higher price somewhat tempers this quantity-increasing effect, but cannot reverse it.

the (actual) consumer surplus. Indeed, the consumer surplus might be negative – the purple triangle might be larger than the red triangle. But whatever consumers lose, the monopolist gains. The purple triangle is part of the blue rectangle, which represents the monopolist's surplus. Therefore, we have a distributional effect, but no reduction in efficiency. Indeed, misperception increases efficiency. By inflating demand, the overestimation bias increases the quantity sold – from Q_M to Q_M' – and thus reduces the monopoly deadweight loss, which is represented by the black triangle. Notice that the black triangle in Figure 2b is smaller than the black triangle in Figure 1b.




Figure 2b: Overestimation, No Price Discrimination, Monopoly

When the misperception is even stronger and the perceived demand curve shifts even higher above the actual demand curve, the quantity, Q_M' , can be larger than Q_C . In this case, the black triangle disappears entirely, and the problem of insufficient purchases is replaced with a problem of excessive purchases. Specifically, consumers in the $[Q_C, Q_M']$ range inefficiently purchase the product. Misperception can either increase or decrease overall efficiency in this market, depending on the relative magnitudes of the insufficient purchases problem (without misperception) and the excessive purchases problem (with misperception).

Finally, consider a monopolistic market with price discrimination. The monopolist charges an individual price for each consumer, based on each consumer's WTP. See Figure 2c. Whereas WTP derived only from preferences in the standard model (see Figure 1c), now WTP is a product of both preferences and misperceptions (see Figure 2c). Price

discrimination allows the monopolist to “march down” the demand curve, setting different prices for different consumers. In the standard model, the monopolist marched down the actual demand curve. Now the monopolist is marching down the perceived demand curve. Turning to welfare, in the standard model the monopolist extracted the entire surplus. Consumers gained nothing, but also lost nothing. Here the monopolist is also extracting perceived surplus, which is represented by the purple trapezoid. This extra gain to the monopolist is a loss to consumers; the purple trapezoid represents a transfer from consumers to the monopolist – a distributional effect with no efficiency implications.⁴⁸ But there are also efficiency implications. Consumers in the $[Q_C, Q_C']$ range should not purchase the product. They buy only because of the misperception, because they overestimate the product’s value. These purchases create an efficiency loss, which is also a loss to consumers (namely, the consumers bear the entire efficiency loss). This loss is represented by the light-red triangle (below the supply curve).⁴⁹

⁴⁸ There are additional distributional implications for the consumer side of the market, especially if we add budget constraints and wealth effects (see supra note 29): Consumers with a high WTP, who would have purchased the product and gained a positive surplus in the absence of price discrimination, lose that positive surplus, and more. Consumers with a low WTP, who would have been priced out of the market in the absence of price discrimination, now purchase the product and pay a price equal to their full WTP, including both the preference-based and misperception-based components. But, as noted above, without misperceptions budget constraints may keep the WTP below the preference-based level; and with misperceptions, if the bias level is sufficiently low, the full WTP may still be below the preference-based level. Therefore, consumers with a low WTP may actually benefit from price discrimination. And if WTP is positively correlated with wealth, then perfect price discrimination may result in progressive redistribution. Still, as compared to the parallel effect in the standard model (see supra note 29), misperception reduces the magnitude of any progressive redistribution.

⁴⁹ The light-red triangle representing the efficiency loss in Figure 2c is identical to the light-red triangle in Figure 2a.




Figure 2c: Overestimation, With Price Discrimination, Monopoly

In the standard model where WTP is derived from preferences alone (Section A above), price discrimination hurts consumers but increases efficiency. Specifically, consumers enjoy no surplus at all, but there is no deadweight loss. With overestimation, price discrimination hurts consumers even more and may either increase or decrease efficiency. Consumers are hurt more, because now they give up surplus that they do not have – perceived surplus – and thus end up with a loss. In terms of efficiency, the insufficient quantity problem is avoided, but an excessive quantity problem is created. Whether price discrimination increases or decreases efficiency depends on the relative magnitudes of the black triangle in Figure 2b and the light-red triangle in Figure 2c. And, when the misperception is stronger such that $Q_{M'}$ is larger than Q_C , price discrimination definitely decreases efficiency. In this case, there is an excessive quantity problem even in the absence of price discrimination, and price discrimination only exacerbates this problem.⁵⁰

⁵⁰ In the standard model, price discrimination has another efficiency advantage: By increasing the provider's surplus, price discrimination increases the availability of products and services that require large fixed cost investments (e.g., a doctor may only move to a small town if she can extract more surplus through price discrimination). And consumers benefit from this increased availability as long as the price discrimination is not perfect. In the behavioral model, even if the price discrimination is not perfect, consumers will likely pay more than the product or service is really worth to them.

III. POLICY IMPLICATIONS

In the standard model, price discrimination is bad for consumers but good for efficiency. The need for legal intervention thus depends on the relative weight of efficiency and distribution in the social objective function.⁵¹ (Unless the adverse distributional implications can be neutralized through an appropriate adjustment to the tax and transfer system.⁵²) Adding demand-inflating misperception, price discrimination is even worse for consumers and might not have any redeeming efficiency advantage. Therefore, when price discrimination targets misperception-based WTP in addition to preference-based WTP the case for legal intervention becomes stronger. In particular, we should consider policies that curb price discrimination, even if we were uncertain about such policies in a world without misperception.

There are different ways to curb algorithmic price discrimination. Policymakers could impose a direct prohibition against price discrimination. But this seems too blunt. They could facilitating market forces that interfere with price discrimination, specifically, by making arbitrage easier and by disseminating information that could trigger a fairness-based consumer backlash.⁵³ (Regulators could also force sellers to disclose the lowest price for which the product or service was recently sold.) Or, policymakers could target the foundation of algorithmic price discrimination: big data. Algorithmic price discrimination is fueled by big data and so this extreme form of price discrimination can be curbed by limiting sellers' access to information about consumers' WTP, namely, by increasing privacy protections and data security measures.⁵⁴ But big data and personalization generate

⁵¹ See, e.g., Steven Shavell, FOUNDATIONS OF ECONOMIC ANALYSIS OF LAW, 597 (2004) (explaining that “[t]here is a vast multitude of ways of aggregating individual utilities into a measure of social welfare,” including some that are purely utilitarian and others that emphasize distributional equity, but that “no single way is endorsed under welfare economics.”).

⁵² See Kaplow and Shavell, *supra* note 6.

⁵³ See *supra* Sec. I.A.

⁵⁴ See, e.g., Federal Trade Commission, “Protecting Consumer Privacy in an Era of Rapid Change,” March 2012 (calling on companies to “build in privacy at every stage of product development,” “give consumers the ability to make decisions about their data,” and “make information collection and use practices transparent”); Natasha Singer, “Why a Push for Online Privacy Is Bogged Down in Washington,” N.Y. TIMES (Feb. 28, 2016), https://www.nytimes.com/2016/02/29/technology/obamas-effort-on-consumer-privacy-falls-short-critics-say.html?_r=0 (describing the Obama Administration’s blueprint Consumer Privacy Bill of Rights and obstacles to its implementation); Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC, 2016 O.J. (L 119) 59 (adopting new data privacy regulations for the European Union); Cedric Burton et al., The Final European Union General Data Protection Regulation, BNA (Feb. 12, 2016), <http://www.bna.com/final-european-union-n57982067329> (summarizing the regulation’s drafting and adoption); Natasha Singer, “Data Protection Laws, An Ocean Apart,” N.Y. TIMES (Feb. 2, 2013), <http://www.nytimes.com/2013/02/03/technology/consumer-data-protection-laws-an-ocean-apart.html> (describing the divergence between recent U.S. and E.U. privacy law reforms). The Do Not Track initiative also holds some promise as a means for stemming the flow of data to sellers who wish to engage in price discrimination. See: <http://donottrack.us>. See also: David C. Vladeck, Digital Marketing, Consumer

benefits as well as costs, and it is not clear whether choking off the flow of information would provide a net gain for consumers.

What about personalized law? Algorithmic price discrimination occurs when sellers gain information about each consumer's individual WTP and set a personalized price equal to that consumer's WTP. When sellers utilize personalized pricing, policymakers can respond with personalized price caps. Price caps are a common policy response to excessive pricing.⁵⁵ The conventional price cap applies uniformly to the entire market. But such a uniform price cap is a very blunt instrument. Why must a rich consumer pay the same, capped price as a poor consumer? Why must a consumer who gains substantial value from a product pay the same price as a consumer who gains much less value from the product? Personalized price caps provide a more subtle response.

The idea of personalized price caps was recently proposed by Ziv Ben-Shahar as a response to fairness concerns about price discrimination in a world without misperception.⁵⁶ But a personalized price cap can also directly address the extra harm that price discrimination inflicts on consumers who overestimate the benefit from a product. Specifically, a price cap equal to the individual consumer's preference-based WTP would neutralize the effects of the misperception. Of course, regulators would need a lot of information to implement such a personalized-price-cap regime. In particular, they would need to distinguish between the preference-based and misperception-based components of each consumer's WTP. But in the world of big data perhaps this obstacle is not insurmountable.⁵⁷

A softer, personalized-law approach replaces price caps with disclosure mandates. While disclosure mandates are a conventional regulatory solution,⁵⁸ this traditional technique has an interesting, non-traditional twist in the present context. In a world of personalized pricing – the world of price discrimination – disclosure mandates should also be personalized. Namely, sellers should be made to provide each consumer with individualized information on the product's true value to that particular consumer. With such a disclosure, consumers will be able to separate out the misperception-based component of their WTP and thus avoid overpaying.⁵⁹ (It is not clear, however, if sellers

Protection, and the First Amendment: A Brief Reply to Professor Ryan Calo, 82 G. Wash. U. L. Rev. (Arguendo) 156, 160-61 (2014).

⁵⁵ See, e.g., Oren Bar-Gill, Price Caps in Multiprice Markets, 44 J. Legal Stud. 453 (2015).

⁵⁶ See Ziv Ben-Shahar, *Ethics and Regulation of Personalized Pricing*, Working Paper (2017).

⁵⁷ Cf. Recent work on personalized default rules, e.g., Cass R. Sunstein, *Empirically Informed Regulation*, 78 U. CHI. L. REV. 1349, 1399 (2011) (“Other default rules are *personalized*, in the sense that they draw on available information about which approach best suits individuals, and potentially even each individual, in the relevant population.”); Ariel Porat & Lior Jacob Strahilevitz, *Personalizing Default Rules and Disclosure with Big Data*, 112 MICH. L. REV. 1417, 1421 (2014) (observing that variable legal defaults may be more effective in changing citizen behavior).

⁵⁸ See OMRI BEN-SHAHAR AND CARL E. SCHNEIDER, *MORE THAN YOU WANTED TO KNOW: THE FAILURE OF MANDATED DISCLOSURE* (Princeton University Press, 2014) (noting the prevalence of disclosure mandates and arguing that they often fail to achieve their goal).

⁵⁹ Regulators around the world have already begun to experiment with personalized disclosure mandates. See OREN BAR-GILL, *SEDUCTION BY CONTRACT*, pp. 32–43 (OUP 2012). See, e.g., *id.* at 37–39 (credit card interest rate and minimum payment disclosures). Compare: Ariel Porat & Lior Jacob Strahilevitz, *Personalizing Default Rules and Disclosure with Big Data*, 112 MICH. L. REV. 1417, 1421 (2014).

would have, or could easily acquire, the information that they are required to disclose. As suggested above, sellers have a strong incentive to identify consumers' overall WTP, but they do not have a strong incentive to disentangle the misperception-based component of the WTP from the preference-based component.)

CONCLUSION

This Article considered the positive and normative implications of sellers' increasing ability to price discriminate – to set personalized prices – when consumers' WTP is the product of both preferences and (mis)perceptions. The shift from a standard, neoclassical model, where WTP reflects only preferences, to a more realistic, behavioral model, where misperceptions significantly influence consumers' WTP, forces a substantial revision of standard results. Price discrimination, which appeared less harmful, if not beneficial, in the neoclassical model, is revealed, in the behavioral model, as a much more damaging, welfare-reducing practice, at least when the misperception results in inflated demand (which, as argued above, is likely the most common case).

Personalized pricing thus requires a policy response. And while existing legal frameworks can be used to combat price discrimination, the most promising or, at least, the most intriguing policy response would fight personalized prices with personalized law. Personalized price caps and personalized disclosure can effectively reduce the adverse effects of price discrimination by curbing sellers' ability to set prices above consumers' true, preference-based WTP. When the market employs personalization, regulators should fight fire with fire and seriously explore the potential of personalized law.

APPENDIX

The appendix considers two extensions. Part A discusses the effects of a positive or negative correlation between the two components of WTP – preferences and misperceptions. Part B studies demand-deflating misperceptions.

A. Misperceptions Correlated with Preference-Based WTP

The analysis in the text assumed that the degree of misperception is not correlated with the consumer's preference-based WTP. Graphically, this assumption was represented by a perceived demand curve that was parallel to the actual demand curve. Put differently, the perceived demand curve was represented by an upward shift from the actual demand curve. What happens when the degree of misperception is correlated with the preference-based WTP?

From an efficiency perspective, the problem with misperception is that marginal consumers with a low preference-based WTP buy a product that they shouldn't buy. Therefore, the efficiency loss depends on the bias level of these marginal consumers. And the bias level of the marginal consumers depends, in turn, on the following two factors: (1) whether the correlation between consumers' bias level and their preference-based WTP is positive or negative; and (2) the per-unit cost of production or, more generally, the location of the supply curve in relation to the demand curve.

The nature of the correlation between the misperception-based and preference-based components of the WTP merits elaboration. When will this correlation be positive and when will it be negative? While a complete typology is beyond the scope of the Article, it is useful to consider the following general scenarios. First, the correlation between consumers' bias levels and their preference-based WTP will be positive, when bias is proportional to (actual) value. For example, assume that consumer A will attend the gym twice a week and thus derive a value of \$1,000 from gym membership; and that consumer B will attend the gym four times a week and thus derive a value of \$2,000 from gym membership. (Consumer A's preference-based WTP is \$1,000; and consumer B's preference-based WTP is \$2,000.) Now assume that both consumers overestimate the frequency of their gym visits by 50% – consumer A thinks that she will attend three times a week and thus derive a value of \$1,500; and consumer B thinks that he will attend six times a week and thus derive a value of \$3,000. The same 50% bias level increases consumer A's WTP by \$500 (= \$1,500 - \$1,000) and consumer B's WTP by \$1,000 (= \$3,000 - \$2,000). Second, the correlation between consumers' bias levels and their preference-based WTP will be negative, when bias is negatively correlated with wealth. It is not that poor people are more prone to bias; rather, rich people can afford to hire expert advisers – human or virtual – that mitigate bias and misperception. And so if preference-

based WTP is positively correlated with wealth and wealth is negatively correlated with bias levels, then the preference-based WTP will be negatively correlated with bias levels.⁶⁰

I start with the positive correlation case. In this case, the distance between the perceived and actual demand curves is larger at the left-hand side of the graph and smaller at the right-hand side of the graph. See Figure 3a. Therefore, when the per-unit production cost is low (i.e., the supply curve is low in relation to the demand curve), the efficiency cost of price discrimination is smaller relative to the uncorrelated misperceptions case.⁶¹ This low per-unit cost case is depicted in Figure 3a. The light-red triangle (below the supply curve) is smaller in Figure 3a, as compared to Figure 2c. In contrast, when the per-unit production cost is high (i.e., the supply curve is high in relation to the demand curve), the efficiency cost of price discrimination is larger relative to the uncorrelated misperceptions case. Turning to the distributional implications of price discrimination, the perceived surplus that the monopolist extracts, represented by the purple trapezoid, is larger relative to the uncorrelated misperceptions case; and this result holds regardless of the magnitude of the per-unit production cost. Intuitively, the positive per-unit cost truncates the perceived demand curve and the overestimation bias; this truncation effect is smaller in the positive correlation case.

⁶⁰ The term “preference-based WTP” may be too narrow, when wealth effects are considered. A more accurate term would be “pre-bias WTP.”

⁶¹ The market for information goods is an example of a market where per-unit production costs are low (or zero).

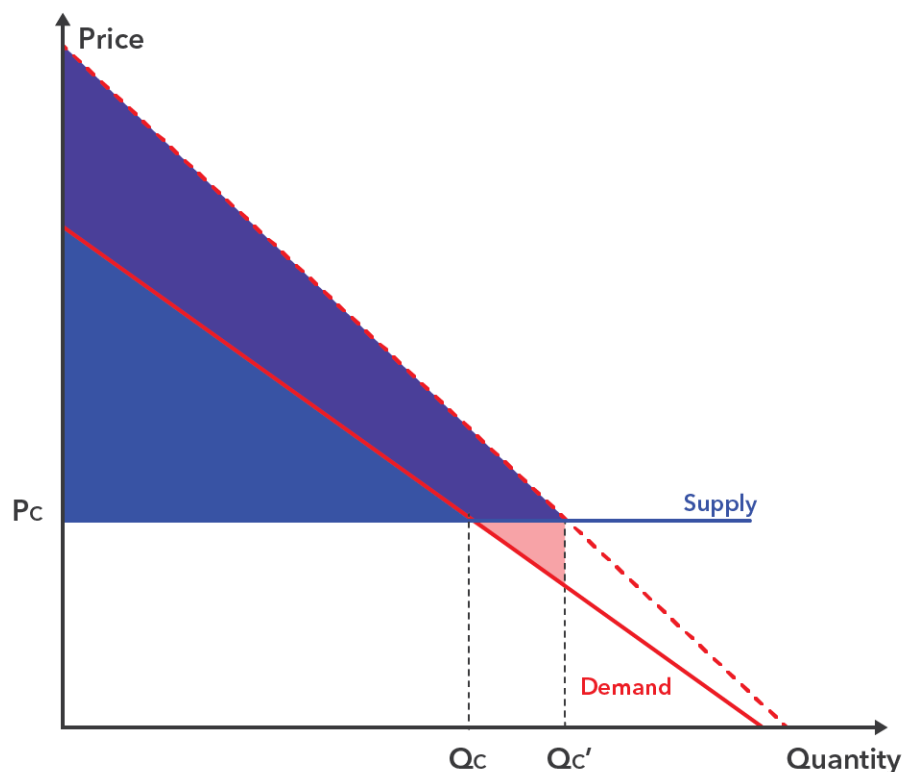


Figure 3a: Correlated Misperceptions - Positive Correlation

Next, consider the negative correlation case. In this case, the distance between the perceived and actual demand curves is smaller at the left-hand side of the graph and larger at the right-hand side of the graph. See Figure 3b. Therefore, when the per-unit production cost is low (i.e., the supply curve is low in relation to the demand curve), the efficiency cost of price discrimination is larger relative to the uncorrelated misperceptions case. This low per-unit cost case is depicted in Figure 3b. The light-red triangle (below the supply curve) is larger in Figure 3b, as compared to Figure 2c. In contrast, when the per-unit production cost is high (i.e., the supply curve is high in relation to the demand curve), the efficiency cost of price discrimination is smaller relative to the uncorrelated misperceptions case. Turning to the distributional implications of price discrimination, the perceived surplus that the monopolist extracts, represented by the purple trapezoid, is smaller relative to the uncorrelated misperceptions case; and this result holds regardless of the magnitude of the per-unit production cost. Intuitively, the positive per-unit cost truncates the perceived demand curve and the overestimation bias; this truncation effect is larger in the negative correlation case.




Figure 3b: Correlated Misperceptions - Negative Correlation

In the standard model, without misperception, price discrimination hurts consumers but increases efficiency. With demand-inflating misperception, price discrimination hurts consumers even more and either increases or decreases efficiency. Correlated misperceptions add nuance to these results. The extra harm that consumers incur is larger in the positive correlation case and smaller in the negative correlation case. In terms of efficiency, price discrimination is more likely to reduce efficiency (1) in the positive correlation case when the per-unit production cost is high, and (2) in the negative correlation case when the per-unit product cost is low. These results are summarized in the following table.

		Per-Unit Production Cost	
		Low	High
Correlation between Bias Level and Preference-based WTP	Positive	Smaller Efficiency Cost Larger Distributional Effect	Larger Efficiency Cost Larger Distributional Effect
	Negative	Larger Efficiency Cost Smaller Distributional Effect	Smaller Efficiency Cost Smaller Distributional Effect

Table 2: Heterogeneous Misperceptions – Welfare Effects

Allowing for a positive or negative correlation between preferences and misperceptions adds nuance to the analysis – either increasing or decreasing the efficiency costs and distributional effects of price discrimination. But the bottom line remains: The behavioral model, where WTP is derived from both preferences and misperceptions, reveals price discrimination to be more harmful than previously believed.

B. Demand-Deflating Misperceptions

The analysis in the text focused on demand-inflating misperceptions, namely, on overestimation of benefit or underestimation of price. I argued that these misperceptions are likely the more prevalent type. Still, in some markets a demand-deflating misperception, either underestimation of value or overestimation of price, can reduce the perceived WTP below the actual WTP. To complete the analysis, I now study such demand-deflating misperceptions and how they interact with algorithmic price discrimination.⁶²

I start, as before, with the perfect competition baseline. The price remains $P_C = k$, but the equilibrium quantity, Q_C' , is smaller. This quantity is determined by the intersection of the supply curve with the perceived demand curve. See Figure 4a. In welfare terms, consumers in the $[0, Q_C']$ range still enjoy the same consumer surplus, as in the no misperception case. This surplus is represented by the red trapezoid. These consumers mistakenly believe that they are getting a smaller surplus – represented by the red triangle – but this false belief does not harm them. There is, however, another group of consumers who are harmed. These consumers, in the $[Q_C', Q_C]$ range, should purchase the product, but don't, because of the misperception. These missed purchases create a welfare loss – an efficiency loss that is born entirely by consumers. This loss is represented by the black triangle.

⁶² The market for health insurance is an example where demand-deflating misperceptions are likely, as consumers underestimate future healthcare costs.




Figure 4a: Underestimation, No Price Discrimination, Perfect Competition

Next, consider a monopolistic market, but without price discrimination. In the standard model, without misperception, the monopoly price is determined by the demand curve. And, when consumers underestimate the value of the product, the price is determined by the perceived demand curve. Therefore, the monopoly price with misperception, P_M' , is lower than the monopoly price without misperception, P_M . The quantity sold with misperception, Q_M' , is also smaller than the quantity sold without misperception, Q_M .⁶³ See Figure 4b. Turning to welfare: In the standard model, without misperception, monopoly pricing prevents some efficient purchases, thus creating the infamous monopoly deadweight loss. The underestimation bias prevents additional, efficient purchases from taking place, thus increasing the deadweight loss. Graphically, we see that the black triangle in Figure 4b is larger than the black triangle in Figure 1b. Note that consumers in the $[0, Q_M']$ range who purchase the product despite the misperception enjoy a larger surplus, thanks to the lower price.

⁶³ The underestimation deflates demand and thus decreases the quantity sold. The lower price somewhat tempers this quantity-decreasing effect, but cannot reverse it.

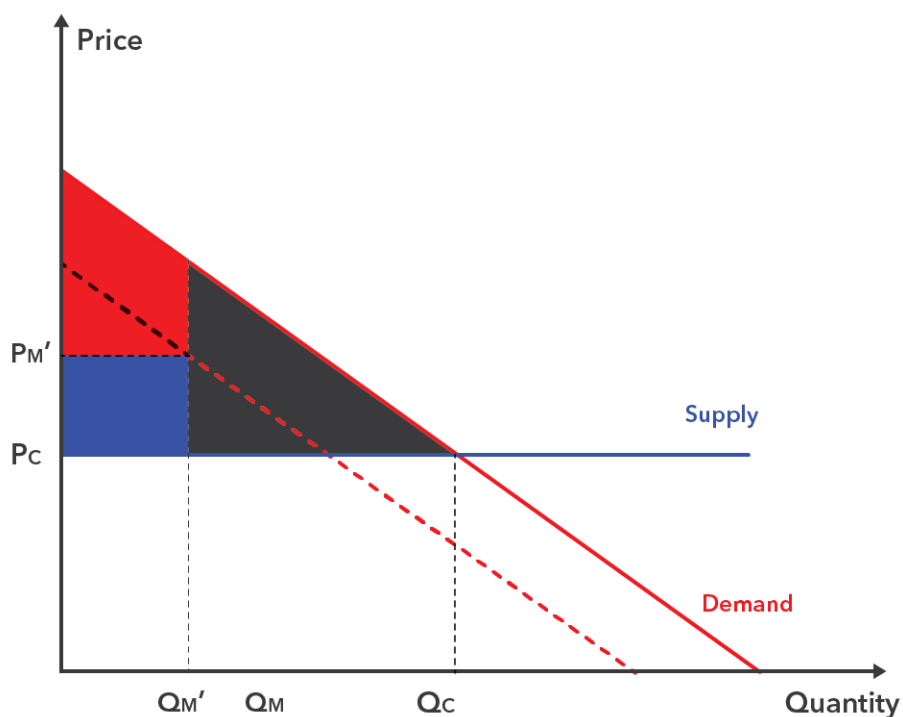


Figure 4b: Underestimation, No Price Discrimination, Monopoly

Finally, consider a monopolistic market with price discrimination. The monopolist charges an individual price for each consumer, based on each consumer's WTP. See Figure 4c. Turning to welfare, price discrimination clearly increases efficiency – it reduces the deadweight loss, i.e., more consumers purchase the product. Graphically, the black triangle in Figure 4c is smaller than the black triangle in Figure 4b.⁶⁴ The effect on the consumer surplus, however, is ambiguous. In the standard model the monopolist extracted the entire surplus. Consumers gained nothing. Here the monopolist can extract only the underestimated perceived surplus, which is represented by the blue triangle. The consumers are left with the difference between the actual surplus and the perceived surplus, which is represented by the red parallelogram. So consumers enjoy a positive surplus, but it is not clear whether this surplus is larger or smaller than the surplus that they enjoy in the absence of price discrimination. On the one hand, more consumers buy the product and enjoy this difference between the actual and perceived surplus. On the other hand, the consumers who would have purchased the product also in the absence of price discrimination enjoy a smaller surplus (because they are charged a higher, personalized price).

⁶⁴ Figure 4c depicts a scenario where $Q_M < Q'_C$. When the underestimation bias is larger, we can also get $Q'_C < Q_M$. But since $Q'_M < Q'_C < Q_C$ in both scenarios, the qualitative comparisons remain unchanged.




Figure 4c: Underestimation, Price Discrimination, Monopoly

In the standard model where WTP is derived from preferences alone (Section II.A. above), price discrimination hurts consumers but increases efficiency. Specifically, consumers enjoy no surplus at all, but there is no deadweight loss. With overestimation, price discrimination hurts consumers even more and may either increase or decrease efficiency. Here, with underestimation, price discrimination clearly increases efficiency and may or may not hurt consumers. These results can be further refined by considering possible correlations – a positive correlation or a negative correlation – between the degree of misperception and consumers’ preference-based WTP (as we did in Section II.C.). But the basic conclusion remains: Algorithmic price discrimination is less worrisome in markets with underestimation. And, as a result, legal intervention is less needed in markets with demand-deflating misperceptions.