Consumer Protection and Disclosure Rules in the Age of Algorithmic Behavior-Based Pricing

Haggai Porat

LL.M. Candidate, 2019/2020, Harvard Law School

LL.M. Paper in conjunction with the Research Seminar in Law, Economics, and Organizations

Supervised by Professor Oren Bar-Gill and Professor Louis Kaplow

24 April 2020

Summary

Two passengers use a ride-hailing service to book an identical trip at the exact same time, but one is charged a higher fare, because a day before she took the bus. Two consumers order the same product from the same online retailer, but one is charged a higher price, because last week he ordered a different product of the more expensive brand. The practice of setting prices based on consumers' prior behavior – as opposed to inherent traits such as race and sex – is called behavior-based pricing (BBP). While many companies are vague as to whether they engage in BBP, the age of algorithmic pricing has placed BBP at the technological frontier, and the profit potential of BBP is quickly becoming irresistible.

The literature on price discrimination has primarily focused on sellers who set different prices based on the inherent features of consumers, such as race and sex. Behavior-based pricing is a fundamentally different practice – prices are set based on consumers' past decisions, which consumers can strategically adjust to affect the prices that they will face in the future. At the same time, sellers can adjust prices to influence consumers' decisions and increase the informational value of these decisions. This paper presents a framework for analyzing the strategic behavior of consumers and sellers in the presence of BBP – both when it is concealed from consumers who are blind to it, and when it is disclosed and made public. Based on this analysis, the paper evaluates the optimal legal response to algorithmic BBP: Should it be banned or regulated? Should sellers be required to make their algorithms transparent or disclose practicing BBP? Should the market be allowed to use BBP freely? The results presented in this paper offer novel, surprising answers to these questions. First, in contrast to the standard economic view that price discrimination is efficient (albeit potentially problematic distributional implications), BBP can either increase or decrease efficiency. Intuitively, customized pricing increases welfare, but obtaining the required information comes at the cost of distorting the decisions of sellers and consumers. Policymakers who are considering the regulation of BBP need to distinguish efficiency-increasing BBP from efficiency-decreasing BBP. The analysis in this paper provides necessary guidance. Second, and counterintuitively, even if BBP is desirable, mandating its disclosure can be harmful to consumers and reduce overall welfare, even if the disclosure does not burden businesses and effectively informs consumers. Intuitively, informed consumers might strategically refrain from purchasing to secure lower prices. This might drive sellers to decrease initial prices. However, it also creates a negative externality on all other consumers in the form of future higher prices, such that requiring businesses to make their algorithms transparent may cause more harm than good.

Recent years have seen a surge in the legal scholarship on algorithmic pricing and how it should be regulated. However, this scholarship tends to bundle together different types of discriminatory pricing practices. This paper shows that the optimal policy response critically depends on the particular type of price discrimination that sellers employ.

Consumer Protection and Disclosure Rules in the Age of Algorithmic Behavior-Based Pricing

Haggai Porat[†]

Abstract

The legal literature on price discrimination has primarily focused on sellers who set different prices based on the inherent features of consumers, such as race and sex. The rise of artificial intelligence and machine learning algorithms has opened the door to a new type of price discrimination: behavior-based pricing (BBP) – the practice of setting prices based on consumers' previous purchasing decisions. Unlike with race- and sex-based discrimination, consumers can strategically adjust these past decisions, to affect the prices that they will face in the future. At the same time, sellers can adjust prices in early periods to influence consumers' decisions and increase the informational value of these decisions, in a way that would maximize sellers' profits in later periods. This paper presents a framework for analyzing the strategic behavior of consumers and sellers in the presence of BBP – both when it is concealed from consumers who are blind to its use, and when it is disclosed and made public. Based on this analysis, the paper evaluates the normative and legal-policy implications of algorithmic BBP. First, in contrast to the standard economic view that price discrimination is efficient albeit with potentially problematic distributional implications, BBP can either increase or decrease efficiency. Intuitively, customized pricing increases welfare, but obtaining the information comes at a potentially higher cost of distorting the decisions of sellers and consumers in earlier stages. Policymakers who are considering the regulation of BBP need to distinguish efficiency-increasing BBP from efficiency-decreasing BBP. Second, and counterintuitively, even if BBP is desirable, mandating its disclosure can be harmful to consumers and might reduce overall welfare, even if the disclosure does not burden businesses and effectively informs consumers. Intuitively, informed consumers may strategically refrain from purchasing to secure lower prices to themselves. This may drive sellers to decrease initial prices, but it may also create a negative externality on other consumers in the form of future higher prices, such that requiring businesses to disclose their algorithms may cause more harm than good.

[†] LL.M. Candidate, Harvard Law School. I thank Professor Oren Bar-Gill and Professor Louis Kaplow for supervising this project, without whom it could not have been completed. For helpful comments and suggestions along the way I thank Meirav Furth-Matzkin, Michael Hornzell, Tamar Kricheli-Katz, Robert Mahari, Ariel Porat, Timna Porat, Steven Shavell, and Tom Zur. All mistakes are my own.

"Dynamic pricing is stupid, because people will find out" (*Spetember 2000, Bill Curry, former Amazon spokesperson, after consumers revealed Amazon's experiment with behavior-based pricing*)[1]

"The promise of dynamic pricing is to increase yield for airlines through targeted offers … We are now making this vision a reality" (*October 2019, a press release by The Airline Tariff Publishing Company*)[2]

1. Introduction

It is well-known that Uber prices surge in rush hours, and that a person that bought a nonstop flight ticket may have paid a higher fare than the person sitting next to her paid for the same flight *and* his next connecting flight combined, just because ordering a nonstop flight demonstrates a higher willingness-to-pay. But does algorithmic pricing stop there? Might it be that ordering an Uber, buying a flight ticket or purchasing a product on Amazon will also increase *future* prices relative to a different person considering an *identical* product (e.g., both ordering a ride or purchasing an identical flight ticket at the same time and place), just because of their prior consumption behavior? The practice of setting prices based on prior behavior of consumers – as opposed to inherent traits such as race or objective market conditions such as high demand during rush hour – is called behavior-based pricing. Amazon experimented with it back in 2000 when it set higher prices to customers that previously purchased certain DVDs, but once consumers found out, Amazon renounced these 'experiments'. More recently, some journalists reported anecdotal concern that Uber may be engaging in behavior-based pricing,[3] but there is no clear proof. However, Uber, as well as other companies, remain vague or silent on the issue. This paper develops a framework that explores how consumers are predicted to change their consumption patterns once they are made aware of behavior-based pricing, and asks whether sellers should be allowed to engage in it, and if they should be obligated to disclose their use of behavior-based pricing? The model presented in this paper offers novel, surprising answers to these questions.

The limited economic literature disagrees whether it is even profitable for a seller to implement this type of behavior-based pricing when consumers are strategically responsive to it. The legal scholarship, on its part, tends to bundle together different types of discriminatory pricing practices, ignoring paramount differences between them. This literature reaches conclusions about regulation of algorithmic pricing that are simply inapplicable to behavior-based, algorithmic pricing. Technological developments in recent

[1] See https://www.washingtonpost.com/archive/politics/2000/09/27/on-the-web-price-tags-blur/14daea51-3a64-488f-8e6b-c1a3654773da/.
[2] See https://www.atpco.net/press/atpco-reduces-barrier-airlines-adopt-dynamic-pricing.
[3] See https://www.theguardian.com/commentisfree/2018/apr/13/uber-lyft-prices-personalized-data.

years, especially the increasing use of neural networks and other types of machine learning algorithms for the pricing of products and services, has made behavior-based pricing inevitable. This paper begins to close the gap between the academic literature and contemporary business practices.

Consumers exposed Amazon for experimenting with behavior-based pricing based on the purchase history of users in as early as 2000. Once discovered, Amazon quickly withdrew it due to public outrage, and hardly anyone followed suit in the next decade. However, significant developments in computer science in recent years, most prominently the commercialization of highly efficient artificial neural network algorithms, have turned this type of behavior-based pricing into a viable, easy to implement, and hard to resist business practice for small and large businesses alike. While historical transactional data is undoubtedly extensively collected by many firms, most remain suspiciously vague as to its precise usage. Most firms vigorously guard their algorithms as trade secrets, pronouncing competitive justifications for the lack of transparency. However, several signs indicate that, below the surface, behavior-based pricing is becoming a central tool that many firms are either developing or already using. Recently, several prominent firms openly admitted to developing tools for behavior-based pricing of this type that they intend to utilize. For example, the Airline Tariff Publishing Company (ATPCO) is co-owned by the three US airline giants and several international airline companies, making it the most significant player in the field of algorithmic pricing of flight tickets. In October 2019, ATPCO announced that it is developing a tool for dynamic pricing, including behavior-based pricing that is based (among other things) on prior transactions of customers.[4] In recent years, there has been a surge in the popularity of off-the-shelf, machine-learning-based pricing software products aimed at non-technical businesses, alongside significant technological strides led by the industry giants. These developments render the penetration of behavior-based pricing inevitable, resurrecting it from the analog age as a new algorithmic pricing frontier.

This paper proposes an analytical framework to formulate an optimal legal response to behavior-based pricing. The results reveal that behavior-based pricing occurs in equilibrium, both when the seller conceals the practice from myopic consumers and when the law mandates disclosure of the practice to consumers that remedies their myopia (who then respond strategically). Initially, all consumers are offered the same price, because they are indistinguishable. However, the decision to purchase the product at a given price conveys valuable information to the seller about the value that different consumers attach to the product, which the seller can use in its subsequent interaction with the consumers. Specifically, the seller offers a higher price to returning customers that previously purchased the product (exhibiting a higher willingness-

[4] See https://www.travelmarketreport.com/articles/Airline-Dynamic-Pricing-Getting-Closer-to-Reality-Says-ATPCO.

to-pay for it) and a lower price to new customers that previously refrained from purchasing the product (exhibiting a lower willingness-to-pay for it). This happens even if consumers are rational, do not develop brand loyalty, and do not learn from using the product, which are the long-standing standard rationales for offering introductory discounts. Indeed, behavior-based pricing is even more general and nuanced, because it can be used to set different prices to two *existing* customers that exhibit different transactional patterns. For example, two customers who purchase a product and each chose a different brand may face a different price in the future for some unrelated product.

Behavior-based pricing presents a unique efficiency trade-off that calls for different legal solutions compared to the paradigmatic practices of price discrimination. One the one hand, behavior-based pricing drives sellers to lower prices, as time progresses, for consumers who exhibit lower willingness-to-pay. This results in more consumers purchasing the product – benefitting both the consumers that can purchase the product at a lower price and the market as a whole. This is the celebrated outcome of most price discrimination schemes, which many scholars and policymakers rely on when arguing against imposing any limitations on price discrimination schemes. However, on the other hand, with behavior-based pricing, there is a cost to getting there. When pricing is based on behavior, rather than inherent traits such as race or sex, the prior decisions of consumers are what generates the informational signal that facilitates future price discrimination. Behavior, by definition, can be controlled, and the prospect of gains (or losses) from price discrimination in the future may distort the incentives of both sellers and consumers in the present.

Sellers, when setting prices, will not consider only their costs and revenue in the present, but also the information that consumers generate when they decide whether or not to purchase, which they can subsequently use to price discriminate. This phenomenon drives sellers to increase prices in the present, which potentially offsets the social benefits from future lower prices. Moreover, if consumers are not myopic or blind to the possibility that a seller may engage in behavior-based pricing, then they may act strategically in anticipation of it and resist the attempt of sellers to learn about their willingness-to-pay for a product or service from their behavior. Specifically, some consumers may find it worthwhile to avoid purchasing a product even if they value it more than its price, forgoing an otherwise efficient deal, to secure a lower price for themselves in the future. Anticipating this, the seller may decide to decrease prices, in order to appeal to some of the strategic consumers. However, such a price decrease may be profitable only to some extent, and the result is that fewer consumers will purchase the product. This dynamic harms both consumers and the market as a whole. Therefore, an initial lesson is that when policymakers and courts consider whether to permit behavior-based pricing techniques, its benefits from subsequent lower prices should be assessed relative to its costs of initial higher prices.

An intermediate legal response is imposing a disclosure mandate, such that businesses that engage in behavior-based pricing will be obligated to make it public. In recent years, mandated transparency has become an extremely popular legal response to the perils of algorithmic discrimination, as advocated by Kleinberg, Ludwig, Mullainathan and Sunstein (2018), Benkler (2019) and others. At first glance, if *effective* disclosure is possible and not too costly, it may seem obviously desirable. When consumers exhibit a limited understanding of some aspect of pricing or marketing, and (1) it is not solved by market forces; (2) disclosure effectively makes consumers rational and informed, rather than takes advantage of their bounded rationality; and (3) the direct costs to disclose are not exceedingly high, then mandating disclosure is commonly perceived as desirable. The vast legal literature on disclosure mandates is focused on extensively scrutinizing each of these assumptions, and algorithmic behavior-based pricing should be subjected to the same scrutiny. However, there is a broad consensus that *if* these assumptions hold, then disclosure is desirable. The basic economic insight is that effective disclosure reduces information asymmetry or consumer bias in the market, which leads to better decisions and superior outcomes. The results presented in this paper challenge this pro-disclosure argument in the context of behavior-based pricing.

Given certain plausible conditions that will be explored in-depth, mandated disclosure of behavior-based pricing is predicted to reduce the overall welfare (as well as the welfare of consumers). One way to frame the intuition of these results is as follows. When consumers are made aware of the fact that a seller engages in behavior-based pricing, some of them react by refraining from a purchase that they would have made otherwise. As previously mentioned, they may do so in order to obtain the benefit of lower future prices that are offered only to those who did not previously purchase the product (or, equivalently, to avoid the penalty of a higher future price). The seller, anticipating this resistance, may decide to lower the price such that it will appeal to more hesitant consumers. However, the outcome is inconclusive and depends on whether this dynamic interaction yields fewer or more consumers that purchase the product in the initial period. Under certain conditions that will be explored, such an effective disclosure, despite improving the decisions of some consumers, may end up harming them as a whole.

To the best of my knowledge, this paper is the first to show that ignorance could be bliss in the context of disclosure to consumers, even if it makes consumers completely rational and there are no costs associated with disclosure or the cognitive decision-making processes of consumers. The source of the possible adverse effect of disclosure lies in a type of collective action problem that has not been previously linked to disclosure mandates. Intuitively, a consumer that strategically refrains from purchasing a product despite its low price, to secure a lower price in the future, is rational. Moreover, a single consumer that would change her behavior would not affect any of the prices. However, the *group* of consumers that strategically refrain from purchasing intentionally become indistinguishable from the group of consumers that did not purchase the product simply because it was too expensive (often, these are the less wealthy consumers). When facing these two indistinguishable groups, the seller is forced to set a single price to everyone in it. Crucially, strategic consumers have higher willingness-to-pay than the consumers who truly do not value the product. This property drives the seller to set a higher price to everyone in this group relative to the price that would have been set if it was not for the strategic group's prior behavior. Therefore, strategic consumers make better choices than myopic consumers, potentially forcing the seller to lower the initial price, but by doing so they also generate a negative externality on their fellow consumers, by driving future prices up. Importantly, this collective action problem is created by the mandated disclosure, and would not have existed if sellers were permitted to remain vague as to their pricing schemes and consumers would remain blind to it.

This is a counterintuitive argument for why, under certain conditions, disclosure and the resulting transparency may result in inefficient outcomes. While courts and policymakers may have in mind other valid justifications for promoting transparency, such as the value that consumers attach to their privacy or the disparate impact of price discrimination on various social groups, the results of this paper demonstrate the importance of tailoring the legal response to the particular properties of any business practice. In the context of algorithmic behavior-based pricing, it should be understood that disclosure and transparency may empower and improve the decisions of some consumers, but potentially at the cost of harming a significantly larger number of consumers, who are typically the less wealthy ones.

Section 2 provides background on behavior-based pricing as a resurging business practice and argues that it is more viable than ever due to recent developments in machine learning, neural network algorithms, and their increasing availability. Section 3 surveys the academic literature on behavior-based pricing. Section 4 uses a formal model to explore how sellers and consumers interact when behavior-based pricing is available and to evaluate the possible legal responses to it. Section 5 discusses how these results should inform legal policy. Section 6 concludes.

2. Background on Behavior-Based Pricing

2.1. Typology

Discriminatory pricing is the practice of setting different prices to different customers (or segments of customers) based on statistically predicted differences in their willingness-to-pay for a product or service or the cost of providing it to them. This includes generally accepted practices such as offering discounts to students and senior citizens in movie theaters, setting higher retail prices in stores located in wealthy neighborhoods, as well as criticized practices such as charging women a different price for roadside assistance or extending loans with higher interest rates to black borrowers based on the probability of default in the group as a whole. Most commonly, price discrimination refers to pricing schemes that set different prices to different people based on their inherent traits (that are given exogenously), such as sex or age. This is a key difference between the pricing schemes that many people have in mind when thinking of price discrimination and other pricing strategies that have increased in prominence in the algorithmic age. Specifically, behavior-based pricing, which conditions prices on the prior behavior of consumers, over which they had control (unlike race and sex).

This common notion of price discrimination based on the inherent traits of consumers is the type that the classic economic literature had in mind when it famously divided price discrimination into three types. Using one or several traits of consumers as indicators of the different willingness-to-pay of each segment is referred to as third-degree price discrimination. In its hypothetical extreme, first-degree (or perfect) price discrimination happens when there is sufficient information to determine the precise willingness-to-pay of every individual consumer. A more nuanced method is second-degree price discrimination, which includes practices such as setting different prices depending on the quantity purchased, or setting different prices for products bundled and sold as a single unit. For example, a luxury car that is offered with an add-on that is priced lower when bought with a cheaper car. These pricing schemes are designed to induce the desired consumption behavior of consumers when the seller has no prior information on the consumers. While this classic taxonomy of discriminatory pricing is often useful, it fails to account for many forms of modern pricing schemes that rely on information that a seller has when setting a price (unlike with second-degree price discrimination), but which is not inherent or exogenous traits of consumers such as race or sex.

A distinct family of price discrimination schemes that are becoming increasingly prominent in the algorithmic age are often referred to as dynamic pricing. This term is somewhat misleading, as some of these pricing strategies are not dynamic in the sense that economists use the term. The grouping together of these pricing strategies under this term may stem from the fact that in practice, a salient feature that they exhibit is that the price offered to the *same* consumer changes across relatively short periods, in addition to the possible variation of prices among different consumers. As a result, there are at least two types of pricing strategies that are often referred to interchangeably as dynamic pricing. One type is surge pricing and the other is behavior-based pricing, which is the focus of this paper.

Surge-pricing is a practice of setting prices based on some current state of the market that changes frequently, but not as a direct result of the actions of any individual person. Most notably, surge pricing happens when sellers adjust prices to the supply and demand that prevail in the market at any given time, such as ride-hailing at rush hour or buying flight tickets at the last minute. Technically, this definition could capture more general economic phenomena, such as changes in prices of products due to an increase of the price of some input required to produce it, although surge pricing is typically not used so broadly. The distinction between surge pricing and "regular" adjustment of prices is not an analytical one, and it seems to vaguely rely on the frequency in which price changes occur.

Note that the different pricing practices that are considered in this section are not mutually exclusive. For example, there are practices that combine regular price discrimination that is based on the inherent traits of consumers with surge pricing. Airline companies often change flight fares not only due to changes in supply and demand but also to target specific consumer segments. For example, there is evidence that prices are higher during the weekend, presumably because it is when people with intensive and high-paying jobs are more likely to find time for it.

Behavior-based pricing is the setting of different prices to the same consumer based on their prior behavior. While behavior is a broad concept that can take many forms, the behavior that is most frequently used for algorithmic behavior-based pricing, and the one that this paper explores, is the transactional history of a consumer in her prior interactions with the same seller. The underlying notion is that consumers vary in the value that they attach to some product or service. Therefore, a decision to purchase a product conveys information as to the willingness-to-pay of the specific consumer.[5] This is

[5] A person's willingness to pay for some product or service is a function of various factors. In the most basic sense, different people may inherently ascribe different values to the same product because of differences in their preferences – the utility that they obtain from consuming the product or service. However, there are other factors that affect willingness-to-pay. Wealthier people are generally willing to pay more for a product, and people that ascribe high value to their time may be decide to purchase a product from the most accessible seller for a higher price if price comparison is time consuming. Notwithstanding the importance of these determinants of willingness-to-pay for empirically assessing it, pricing decisions generally relies on the single aggregated measure. For this reason, this paper simplifies away from questions that pertain to the reason that someone is more or less willing to pay for a product, by focusing on this aggregated measure and not on its components.

distinct from a related notion that purchasing a product can induce actual changes in the willingness-to-pay of consumers, either because they may develop brand loyalty or simply learn about the value of the product through experiencing it. Moreover, the behavior consumers may convey other types of information, such as the costs of providing the product to that specific consumer. For example, when a consumer repeatedly returns items to the store or makes a late payment on an installment loan. While this paper offers insights that may further the understanding of various types of behavior-based pricing, the model explicitly considers only the case where pricing is based on prior purchasing decisions of consumers.

While behavior-based pricing is nowadays closely related to algorithmic pricing, its origins are significantly older. For example, offering discounts to new customers is a type of behavior-based pricing, since it conditions prices on past purchases, albeit in a rather simplistic fashion. The prevalent explanation in the literature for introductory pricing is that consumers learn about the value of the product after experiencing it, such that if consumers underestimate the value of a product before trying it then lower introductory can be used to attract these consumers. However, it could also be justified if there is no such learning and consumers possess accurate information on the value of the product beforehand. The reason is that existing customers have already demonstrated that they value the product more than others that decided not to purchase it – buying the product is a signal of the willingness-to-pay of a consumer. Admittedly, this is a very crude measure, as it is based on a single action (whether a person purchased in the past or not), and the signal is possibly very noisy. For example, a person may purchase a product for the first time not because they value it less than people who purchase it twice, but simply because the need for it did not arise earlier. For these and other reasons, while behavior-based pricing was known to economists, it did not feature as a prominent business strategy relative to other forms of price discrimination, such as race and sex based discrimination.

However, the rise of online markets and algorithmic pricing have dramatically increased the potential profits that could be made from this type of pricing. For example, ride-hailing services could base their prices on the previous rides that the user ordered, and online retailers can price products based on the customer's purchase history. At the same time, there are promising signs that many prominent types of discrimination based on inherent features, such as race as sex, are gradually decreasing over the last decade (although, unfortunately, these problems are unlikely to disappear in the foreseeable future). To understand the resurgence of behavior-based pricing in recent years, which will be described in the next section, it is important to understand the conditions that facilitate it and increases the potential gains from it.

In a nutshell, behavior-based pricing has the greatest potential in markets where consumers repeatedly interact with the same seller, and when the seller can identify consumers, set different prices, and collect transactional data to be matched to the consumer when subsequent interactions occur. While many consumers repeatedly interact with Walmart in their brick-and-mortar stores, it is clearly in a bad position to engage in behavior-based pricing. It is difficult for Walmart to set different prices to different consumers (possibly through personal discounts), and it is also difficult to keep track of the identity of consumers. The example of Walmart demonstrates, by negation, the crucial role of online markets in the resurrection of behavior-based pricing. The conditions listed above are easily met by airline companies, retailers such as Amazon and Alibaba, ride-hailing services, and many more. However, it is important to keep in mind that behavior-based pricing can be viable even if some of the conditions are not met. For example, even though borrowers interact with different lenders, the credit score system overcomes this obstacle by making public information on the prior behavior of borrowers.

As alluded to earlier, a central property of behavior-based pricing that drives much of the results presented in this paper is that behavior – unlike inherent traits such as sex and age – is determined by the decisions of consumers. This fundamental difference is the source of the concern that consumers may distort their consumption behavior in response to dynamic pricing, in anticipation of how their behavior will affect future pricing. The strategic interaction between consumers and sellers does not exist when price discrimination is based on some preexisting exogenous traits such as age, sex or race. Nor does it exist with surge pricing, which is triggered by market forces that are, by definition, outside the control of any single consumer.

To be clear, there are some things that consumers may do to affect pricing in response to price discrimination that is based on inherent traits and surge pricing. Consumers can try to conceal or misrepresent their age or race to avoid higher prices, or to engage in arbitrage by purchasing a product for the lower price and sell it in a secondary market to consumers who were offered a higher price. However, it is generally difficult to conceal these features, especially in contexts where initial registration is required to receive a price offer, and it is only becoming increasingly more difficult with the use of image recognition technologies and various "tells" that machine learning algorithm easily pick even when the information is not explicitly disclosed or verified. Moreover, if a consumer misrepresents her age to get a lower price, it may very well be for the better, as it may result in an efficient purchase of a product that otherwise may not have been purchased.[6] As to arbitrage, its viability is limited by transaction costs, contractual obligations, and registration requirements. As to surge pricing, consumers can respond to these pricing schemes by merely waiting for market conditions to change. Often, however, this is precisely the desired outcome – when rush hour demand for rides exceeds the supply of drivers, it is efficient that those who value the ride more will pay more for it, while others use public transportation or wait for the surge to pass, if possible.

The possibility of strategic resistance of consumers to price discrimination is a unique and dominant feature of behavior-based pricing. While economists generally consider price discrimination based on inherent traits and surge-pricing to be efficient, this fundamental property of price discrimination may not hold for behavior-based pricing.[7] Clearly, there is an empirical question of whether consumers are actually reactive to any sort of discriminatory pricing. Li et al. (2014) found that in 2005 it was already the case that 5-20% of consumers in the air-travel industry exhibited strategic behavior by timing their purchases to periods when prices decrease. Moreover, for consumers to react strategically to pricing strategies, it is often not necessary that they engage in any conscious deliberation, which many time-constrained and rationally-bounded consumers are unlikely to do. Algorithms are gradually being developed not only for sellers but also for consumers, such that actual awareness of consumers is often not required to induce a strategic response. For example, several flight search engines (e.g., Skiplagged) allow consumers to search for flights in a way that circumvents the higher price that the algorithm of the airlines would set if the consumer approached it directly.[8] In the ride-hailing market, Bellhop provides an automated price comparison service between Uber and Lyft, which could potentially account for the firms' pricing strategies.

[6] "Hidden-cities ticketing" provides a recent illustration of the potential to effectively fight static price discrimination. It was recently exposed that many airline companies often charge a higher fare for a direct flight from A to B than a flight from A to C that stops in B, even though the first leg is identical, because people who purchase direct flights have a sufficiently higher willingness-to-pay than those who purchase connecting flights. "Hidden-cities ticketing" is the practice of buying the cheaper ticket, from A to C, and staying at B upon arrival, leaving the part of the ticket from B to C unused. Several flight search engines already incorporated this insight in their algorithms, to the discontent of airline companies. While not using a ticket creates a social waste in the short run, barring airline companies from this practice will efficiently reduce the supra-competitive prices of nonstop flights (although possibly at the cost of higher prices for connecting flights). See https://www.economist.com/gulliver/2015/01/12/airlines-to-public-please-ignore-this-blog-post.

[7] There are, however, other important sources of inefficiencies of price discrimination, most notably from the behavioral literature. For example, Bar-Gill (2019) and Bar-Gill (2020) find, in different settings, that price discrimination may be inefficient if consumers misperceive the value of a product.

[8] See footnote 6, *supra* for a detailed description of this practice.

The final feature of behavior-based pricing that drives some of the results of this paper, is that sellers may find it challenging to commit to not using behavior-based pricing, even if they wanted to. A seller could declare that it does not use past purchases to set prices, but it is unclear whether the seller will resist the future temptation of breaking this promise, especially since this type of information is collected anyway, for various legitimate reasons. Surely, the promise could be credible if sellers value their reputation and behavior-based pricing is easily detectible, but this is often not the case. Many prominent products and services experience price changes for many reasons, and it could be exceedingly difficult to determine whether a change in price resulted from behavior-based pricing. For example, an extreme example is ride-hailing services, where prices change from minute to minute, and from one block to the next. Online retail sellers like Amazon are probably easier to detect, but even then price could vary based on the time of day or place of residence, which – while possibly concerning – are different from behavior-based pricing in important ways. Therefore, there are many cases where only a controlled study could be used to scrutinize the pricing algorithm and detect behavior-based pricing with a sufficiently high degree of certainty.

2.2. The Rise of Behavior-Based Pricing

Throughout the 20th century and the first decade of the 21st century, behavior-based pricing was applicable in a rather narrow set of circumstances. One of the prominent examples is the widespread but rather specific practice of discounts to new customers, rendering behavior-based pricing less significant when compared to the widespread use of discriminatory pricing based on inherent traits and surge pricing. Therefore, it is not surprising that most of the legal literature wholly overlooks the distinctions between these pricing practices. However, this focus is largely outdated, with the last couple of years marking the turn of the tide. In recent years, we have seen an exponential increase in the availability and usage of online web- and application-based marketplaces. From buying products such as groceries, clothing, furniture, and cars, to acquiring services such as ride-hailing, home and car rental, loans, laundry, food delivery and even medical consulting and higher education, there are hardly any consumer markets that are not either dominated or significantly penetrated by online providers.

An important consequence of this revolution is the ability of sellers to observe, record and store voluminous data on the activity of consumers in extreme detail, which nearly all of them already do. The ease of tracking and recording the behavior of consumers dramatically increases the potential profits from behavior-based pricing, and extends its viability to markets where it was not previously thought to be applicable. Another important feature is that many new markets are occupied by few players, inducing a tendency of consumers to repeatedly interact with the same seller, which further facilitates behavior-based pricing. The reality in many prominent markets has fallen short of the early promise of online markets to promote competition through lower entry barriers, easier price comparison and the elimination of geographic monopolies.

Furthermore, technological developments significantly increased the accuracy and complexity in which behavior-based pricing can be performed. First, the increasing computational power and data transfer speed make it possible to engage in complex algorithmic predictions and pricing at a speed that meets the high standards of commercial use. More importantly, recent years have brought groundbreaking developments in machine learning, specifically artificial neural network algorithms. These tools, that giants like Google and Uber continuously work to improve for their usage, have had a profound impact on the ability to translate the economic idea of behavior-based pricing to highly efficient algorithms. While state-of-the-art algorithms are often kept secret, there are many companies that offer efficient and accessible software that any small business can purchase at a low cost, making behavior-based pricing widely available.

Perhaps the first to experiment with online behavior-based pricing was Amazon, as early as in the year 2000, by setting higher prices to consumers that previously purchased certain DVDs. However, it was quickly revealed and abolished due to wide public criticism.[9] Bill Curry, then Amazon spokesperson, interestingly said that "Dynamic pricing is stupid, because people will find out."[10] Possibly due to Amazon's negative experience, alongside the absence of adequate algorithms at the time, it seems that for many years dynamic pricing was either not widely practiced or else kept better hidden. The Amazon anecdote is a useful benchmark to examine the shift in the industry in the following two decades, best demonstrated by the recent case studies of Uber and The Airline Tariff Publishing Company (ATPCO).

Uber is the world's largest company offering ride-hailing services by, in essence, connecting users to a vast network of drivers. Uber determines the fare upfront as a lump sum that is presented to the user as a take-it-or-leave-it offer. Initially, the fare was thought to be determined mostly based on information such as the driving distance and driving time as predicted by Uber's algorithm. At some point Uber started to use surge pricing, adjusting prices to the supply and demand at any given time and geographical area. Chen, Mislove, and Wilson (2015) found out that Uber uses surge pricing extensively before it was publicly acknowledged. Unlike Amazon, Uber's response was not to forgo the practice but rather to make it visible

[9] See http://edition.cnn.com/2005/LAW/06/24/ramasastry.website.prices/.
[10] See https://www.washingtonpost.com/archive/politics/2000/09/27/on-the-web-price-tags-blur/14daea51-3a64-488f-8e6b-c1a3654773da/.

to the user, disclose it in its terms and conditions and attempt to sway the public opinion in favor of it. In doing so, Uber emphasizes the efficiency justifications of a pricing scheme that is attentive to demand and supply, which benefits consumers as well as drivers.[11] In contrast to its extensive efforts to sell the idea of surge pricing, it remains silent on whether it also engages in behavior-based pricing.

The anecdotes of Amazon and Uber raise the question of whether Uber, Lyft or other similar companies that use algorithmic pricing, may also be engaging in concealed behavior-based pricing. In other words, whether they use the transactional history of users to predict their willingness-to-pay for the services or products offered and set the prices accordingly. For example, Lyft may benefit from setting a higher price to a person who ordered a ride during rush hour in the past, instead of taking the bus; or a lower price to a person who declined the same price offer in the past, and took the bus or rode her bicycle instead (where declining could be inferred, for example, from opening the application and closing it without ordering a ride, from opening it for some minimal time, from soliciting a price after entering the desired destination, etc.). In a more straightforward fashion, Alibaba or Amazon may benefit from setting a higher price for practically any product for someone who previously purchased some luxury jewelry, simply because they are more likely to be wealthy.

Some, mostly technically-oriented bloggers but also journalists, take it for granted that data on prior behavior is probably used to set prices.[12] Surprisingly, though, it seems that this issue has mostly escaped the public's eye. Generally, the public seems to be highly aware that algorithms are used to extract higher prices from consumers based on various information, with one survey estimating that 88% of Chinese consumers are suspicious.[13] However, often these are vague concerns that do not necessarily relate to any specific type of information collected or method used to set prices. To the best of my knowledge, Amazon, Alibaba, Uber, etc. have yet to be publicly confronted with the specific question of behavior-based pricing based on transactional history.

If Alibaba or Amazon were engaging in behavior-based pricing, it would seem reasonable that they would hide doing so given the possible legal and reputational consequences. When justifying surge pricing, Uber often explains that rush hour or lousy weather makes it impossible to provide the service for everyone demanding it, implicitly conveying the notion that it is a necessary evil rather than a means of increasing Uber's profits. In contrast, using a person's past inclination to pay high prices as a standalone factor to set

[11] See https://marketplace.uber.com/pricing/surge-pricing.
[12] See https://www.theguardian.com/commentisfree/2018/apr/13/uber-lyft-prices-personalized-data and https://datarootlabs.com/blog/uber-lift-gett-surge-pricing-algorithms.
[13] See https://technode.com/2019/03/28/beijing-consumers-association-price-discrimination/.

a higher price in the present may trigger harsh criticism, as happened with Amazon. However, behavior-based pricing systematically may also result, indirectly, in lower prices to groups in society that have a lower-than-average willingness-to-pay. For example, a person of color with less-than-average income might pay a lower price based on her purchasing decisions, signaling her willingness-to-pay, even if it is impossible to observe her race or ethnicity directly.

Notwithstanding the force of the reputational concerns, the model in this paper highlights another source of concern, that explains why sellers may want consumers to be aware of surge pricing, but not of behavior-based pricing. Surge pricing is aimed at aligning supply with demand, by encouraging consumers to time their purchase based on prices. In contrast, the response of consumers to behavior-based pricing may be detrimental to the profits of sellers, as some consumers may resist the attempt of the seller to collect reliable information by avoiding purchases and disguising to enjoy lower prices in the future. Therefore, even if there are no reputational forces, there is a structural reason to believe that sellers may be more determined to conceal behavior-based pricing. This does not mean, however, that the law should mandate seller to reveal it. This paper reveals that indeed, and counterintuitively, it may be better to allow sellers to keep it hidden.

Currently, the law does not mandate the disclosure of any type of information regarding algorithmic pricing, and most companies choose not to disclose how their algorithms work with sufficient detail that would allow determining whether behavior-based pricing is practiced. However, the law does require sellers to disclose certain things that may provide hints. Most importantly, sellers must specify in their terms and conditions what data they collect on users, as well as some (possibly vague) description of its usage.

Taking Uber as a case study, it seems that it does not explicitly deny using the prior activity of users as input in its pricing algorithm, vaguely stating that "driver demand" is used to set the price, and that "Our understanding of a city improves over time as more people take more trips across different routes. We use this understanding to make better predictions about when and where riders will request rides."[14] A careful reading of the Uber Privacy Notice[15] further reveals that Uber collects and stores practically all historical user data. This includes all information relating to a ride that was requested, but also when the user elicited a price offer for a ride but decided to decline it by not ordering the ride (possibly because the price was too high). Furthermore, Uber collects data on the location of users while using the application,

[14] See https://marketplace.uber.com/pricing/route-based-pricing.
[15] See https://www.uber.com/global/en/privacy/notice/.

which admittedly facilitates several features in the user interface. Oddly, however, for users outside the EU (who do not enjoy its strict privacy protection), Uber also collects the location of the user when the application is not in use but left open in the background (either by accident or from indifference to closing it). It is unclear what functional purposes could be served by collecting this data. However, this information could be used to easily infer what a user decided to do after declining a price offer that was elicited, generating valuable information as to her willingness to pay for a ride. She might have walked to her destination, cycled, driven by some other means, perhaps along the bus route, implying that she took the bus instead, or perhaps she stayed put. This rich data seems to be highly predictive of the willingness-to-pay of users that declined rides. Uber's privacy notice further states that Uber uses the collected personal data to automatically set the price of a ride, "based on constantly varying factors," which is defined as an open list ("such as…"), as well as to "personalize" promotional offers (often in the form of discounts), for which it uses – here less vaguely – all data "including … past use of Uber's services."

These hints offer a possible solution to a puzzling feature of the service offered by Uber. Uber's interface conveniently includes the option of taking the bus – one of Uber's most significant competitors. The option appears just below the option to request a ride, and upon selection displays features such as suggesting the best bus route to take to get as fast as possible to the destination, the real-time location and estimated time of arrival of the next bus (to decrease waiting time), and more. Indeed, it may be profitable for a well-established service provider to supply a wholesome service and keep users happy even at the cost of facilitating occasional transactions between its users and its competitors. Moreover, since busses are slower, it may sometimes encourage customers to choose an Uber ride instead. Not dismissing these possible motivations, it is clear that this feature also tells Uber what a user did after rejecting a price offer, which is valuable information if Uber wanted to engage in behavior-based pricing for future rides.

The secrecy and vagueness surrounding numerous companies that use algorithms to set prices to consumers persist to-date. Suspicious consumers quickly detected Amazon's experimenting with behavior-based pricing in 2000, but this is unlikely to happen today. With Amazon, all that was required was for a user to delete her cookies[16] to observe a price change. Nowadays, however, nearly all sellers require users to log in to a registered account to initiate a transaction, making cookies and anonymous purchases somewhat obsolete (as well as arbitrage, in some cases). Moreover, consumers came to accept that prices change for a variety of reasons, often from minute to minute. Therefore, it is not trivial to infer whether a specific price change is a result of behavior-based pricing based on prior activity or whether it

[16] Files saved locally on the computer of the user that allows a website to identify the user in the future.

is a result of any number of other factors. However, these practices are susceptible to methodical experimental field testing, applying more rigorous identification strategies. Such studies include Cui et al. (2018), that found price discrimination by Alibaba, and Chen, Mislove and Wilson (2015), that found surge pricing by Uber before it was publicly acknowledged. To the best of my knowledge, no similar attempt to explore behavior-based pricing experimentally was published or otherwise made public.

Prominent players that could adopt behavior-based pricing practically immediately without being easily detected, if they have not done so already, include giants such as Amazon (that admitted to doing it in the past), Uber, Lyft, Asos, Alibaba, and many others. However, while the potential implications of this largely-overlooked pricing strategy cannot be overstated, the vague cases are not the only ones out there. There are several prominent examples where behavior-based pricing based on prior activity is considered best practice, unhidden and often prided.

The Airline Tariff Publishing Company (ATPCO) is a privately held corporation co-owned by most of the world's largest airline companies, including United Airlines, Delta, FedEx, American Airlines, Lufthansa and many others. ATPCO is dedicated to collecting and distributing fare-related data as well as to developing data-driven algorithms for pricing flight fares.[17] In October 2019 ATPCO announced that it is getting closer to developing a tool to implement "dynamic" pricing based on, among other factors, the consumer's flight history.[18] In its press release, ATPCO stated that "The promise of dynamic pricing is to increase yield for airlines through targeted offers that meet consumer expectations … We are now making this vision a reality."[19] The surprisingly candid response of Aldo Ponticelli, head of distribution strategy at Alitalia, was that "The journey to personalization has begun, but it's a marathon, not a sprint, and we are just at the starting line."[20]

The state-of-the-art algorithms are developed by industry giants, like Uber, Google, and Netflix, who tailor them to their specific needs. However, many software companies develop standardized algorithms for behavior-based pricing that are offered for purchase as friendly and accessible "business management" licensed software. With prices as low as $500 a month, these software programs are available for any small business that wishes to engage in algorithmic pricing, including behavior-based pricing.

---

[17] This unique cooperation between airline companies drove the US Department of Justice to launch an investigation into ATPCO on antitrust grounds, but eventually the case settled and ATPCO continued to operate (See Borenstein 1992).
[18] See https://www.travelmarketreport.com/articles/Airline-Dynamic-Pricing-Getting-Closer-to-Reality-Says-ATPCO.
[19] See https://www.atpco.net/press/atpco-reduces-barrier-airlines-adopt-dynamic-pricing.
[20] *Id*.

Hubspot is a software company headquartered in Cambridge, Massachusetts. At about $8 billion in market value and $500 million in annual revenue, Hubspot is one of the leading software companies that develops and customizes sales and marketing software tools for 70,000 businesses worldwide, among them the auto manufacturer Subaru and the popular food delivery service Doordash. For obvious reasons, Hubspot takes pride in its clientele but does not supply detailed information on the specific tools sought by any single client. However, for only $500 a month, any small business can purchase a "Sales Hub," which is a user-friendly software for automatization of sales strategy. One prominent tool that is included in the Sales Hub is "predictive lead[21] scoring," a tool to measure the likelihood that a potential customer will make a purchase. The score is "based on any characteristic or behavior … with your historical data in mind,"[22] with categories of data such as 'Email engagement,' 'Demographic information,' and 'Behavior.' Clients can further enhance this strategy in many ways, for example, by integrating a software developed by Hubspot's partner Wistia (specializing in videos) that allows "tracking [the lead's] viewing session second by second … data can be used to update lead scores."[23]

PerfectPrice is a boutique software firm founded in 2014 that specializes in developing tools for all types of discriminatory pricing. As to behavior-based pricing based on prior user activity, PerfectPrice's promotional video states that "The AI's understanding of your customers' behavior improves over time, becoming a dependable and accurate way for you to deploy prices."[24] Interestingly, PerfectPrice is a business partner of ATPCO, and its prominent advisor is Robert Phillips, former director of Uber's Marketplace Optimization department and current director of Amazon's Pricing Research department. In 2012 Phillips, then a Columbia University Professor of Professional Practice, co-edited The Oxford Handbook of Pricing Management, contributing an article on Customized Pricing that states that "Relative to its [i.e., customized pricing] importance in the economy, it has been relatively less studied than other pricing modalities ... There are at least two promising areas for future research. One is the estimation of customer price-sensitivity information from historical data."[25]

To summarize, behavior-based pricing based on prior consumer activity is deep-rooted in several pre-internet markets. However, the overwhelming penetration of online markets and recent developments in computer science brought a resurgence of behavior-based pricing in greater scope and potential for

[21] The term in the marketing jargon for a potential customer contemplating whether to make a transaction.
[22] See https://www.hubspot.com/products/marketing/lead-management.
[23] See https://ecosystem.hubspot.com/marketplace/apps/marketing/video/wistia.
[24] See https://www.perfectprice.com/.
[25] Phillips (2012, 28).

businesses than ever before, increasing the gap between the academic literature and contemporary business practices and technological innovation. All indications suggest that this trend is likely to continue and intensify, for better or worse, as long as no legal intervention is administered. The question remains, however, what type of intervention is desirable – should behavior-based pricing be banned? Should sellers that practice it be obligated to disclose this fact, or neither? Building on the academic literature that will be surveyed in the next section, the model presented in this paper provides some surprising answers to these questions.

3. Related Literature

3.1. Behavior-Based Pricing

Rothschild (1974) was the first to argue that sellers can use dynamic pricing not only to respond to changes in demand, which I refer to as surge pricing, but also to reveal and learn the aggregate demand function when it is initially unknown to the seller. However, this and many subsequent studies either focused on learning the aggregate demand function, rather than the demand of any specific consumer, or assumed that consumers do not react to anticipated pricing decisions.

The first study that accounted for the fact that buyers may alter their behavior to affect the information that the seller receives from their behavior is Hart and Tirole (1988). While this was not the primary purpose of the study, it explores, among others, a setting where a single seller can lease a good to a single buyer repeatedly, without initially knowing whether the buyer ascribes a high or a low value to the good, but potentially learning. They find, however, that a seller cannot profit from setting different prices at different periods. This result is primarily driven by the fact that a high-valuing buyer understands that their decision may reveal their type and affect future pricing. Given the specific setting of this model, the buyer always acts in a way that conceals their type, and therefore in equilibrium, the seller will not learn anything about the buyer's type.

Nowadays sellers use consumer data both to increase the value of products to consumers (through personalization of the product, among other means) and, often concurrently, to extract additional profits (through price discrimination, for example). In the first half of the internet age, however, gathering consumer data was mostly done for the latter purpose, and therefore practices such as selling customers lists and storing browsing 'cookies' were seen as purely harmful for consumers. Therefore, despite the fact that consumer privacy and distorting consumption behavior are distinct issues, it is not surprising that the first studies to explore behavior-based pricing when consumers are strategic emerged from the privacy literature.

Taylor (2004) considers a framework where the prior transactions of consumers are observed through the trading of customer lists between sellers, but it directly applies also to cases where consumers repeatedly interact with the same seller. Taylor finds that if there are two types of strategic customers, with either low or high valuation for a product, then some of the high-valuing consumers will strategically avoid an efficient purchase to prevent the second seller from using this information to set a higher price for them. Strikingly, however, as in Hart and Tirole (1988), customers in equilibrium are able to conceal their type and sellers fail to learn anything. Importantly, though, these results were obtained in a setting where there are only two types of consumers, which allows the seller in the second period to extract the entire surplus from customers that are identified as having high valuation, creating a powerful incentive to conceal. The model in this paper considers the more realistic assumption that consumers have varying levels of valuation for the product, inducing behavior-based pricing in equilibrium, because only some of the consumers conceal themselves.

Acquisti and Varian (2005) emphasize that the theoretical finding that sellers cannot profitably engage in behavior-based pricing and learn the valuation of consumers is puzzling, given the fact that sellers are known to invest heavily in gathering consumer information. In their setting, consumers decide whether to purchase, but importantly, they can also directly erase their data, analogous to the ability to delete 'cookies'. However, their main results are similar, that a monopolistic seller is unable to engage in behavior-based pricing in equilibrium, and therefore consumers do not bother erasing their data (although the threat of doing so is important, as it may induce the seller to refrain from engaging in dynamic pricing). They find that behavior-based pricing based on purchase history is profitable only if some of the consumers are myopic, do not anticipate the future behavior-based pricing, or if a purchase in the second period creates some additional value for those who previously purchased the product – some added value from having two products. The similar and surprising result of this study may also be attributed to the fact that it assumed that there are only two types of consumers, of either low or high valuation of the product. In their model, a high-value type that reveals herself as such expects to receive zero value from the future transaction. The intuition is that absent a credible commitment by the seller to set future prices, it is expected to set the highest price possible that extracts all consumer surplus from perfectly-identified consumers. Since there are only two types of consumers, any separation in equilibrium perfectly identifies those consumers that are of high-value and purchase the product.

Fudenberg and Villas-Boas (2006) is the first study, to my knowledge, to assume that there are more than two types of consumers where consumers behave strategically and, consequently, find that behavior-based pricing occurs in equilibrium. When there are many consumers that value the product differently, only some of them decide to forgo the purchase in the first period in order to obtain the future lower price, resulting in a separating equilibrium. In this setting, the price in the first period spearates the market into two segments. One segment includes the consumers that value the product less than its price, as well as consumers that value the product only somewhat higher than its price, who are the strategic consumers that refrain from purchasing the product to secure a future lower price. The other segment are consumers who sufficiently value the product to decide to purchase it despite anticipating the higher price in the future. Consequently, in the second period, each segment is offered a different price, where those that previously purchased the product are offered a higher price and those that did not previously purchase the product are offered a lower price. Building on this setting, the model in this paper explores the additional possibility that behavior-based pricing is concealed by the seller and asks whether mandating disclosure of behavior-based pricing is desirable.

In contrast to previous studies, Chen and Zhang (2009) show that even competing firms can benefit from behavior-based pricing based on purchasing history when consumers act strategically, but only if some consumers are loyal to a brand and others are 'switchers'. It turns out that in this setting behavior-based pricing is profitable due to an anti-competitive effect that is facilitated by loyal customers. In equilibrium, however, no consumer has an incentive to alter her purchasing behavior to affect future pricing. This is somewhat counter-intuitive considering the realistic settings where purchasing does in fact trigger higher future prices.

The literature surveyed so far focused on behavior-based pricing based on purchase history. Another closely related type of consumer behavior that could be distorted by behavior-based pricing is the timing of the purchase. Su (2007) and Liu and van Ryzin (2008) show that consumers may profit from postponing an efficient purchase to enjoy a lower future price. However, this is limited to a setting where the seller commits to a future pricing schedule of the product, while most of the prior literature, as well as this paper, assumes that sellers cannot credibly commit to strategies that are not sub-game perfect. Li et al. (2014) show empirically that 5-20% of consumers in the air-travel industry are strategic in the sense of timing purchases when the price decreases.

3.2. Algorithmic Transparency

The notion of protecting consumers by imposing legal limitations on the use of algorithms by sellers is a relatively new debate that has experienced a surge in recent years. Approaches to regulating algorithms include banning the deployment of algorithms for specific purposes, as well as imposing restrictions on the input that may be fed to an algorithm (See Kleinberg et al., 2018; Yang and Dobbie, work in progress). Another legal tool that is probably the most debated in recent years is that of transparency – imposing duties on actors in the private sector to make their algorithms public. While there are important differences between making an entire code of an algorithm public and other levels of transparency, such as a duty to disclose some of the specific functions that the algorithm performs, transparency is mostly an equivalent term to mandated disclosure – both are about forcing sellers to reveal private information regarding their products or business practices (and are, therefore, used interchangeably).

Given the technological properties of algorithms, it is not surprising that transparency received the most attention. As stated in Benkler (2019), "Inside an algorithmic black box, societal biases are rendered invisible and unaccountable." With the backdrop of the technological properties of algorithms, Kroll et al. (2017) argue that transparency is not sufficient for regulating algorithms. Instead, they provide a set of technological tools that can be used to test the compliance of algorithms with legal anti-discrimination standards, such as avoiding discriminatory training data, maintaining randomness, manipulating data in favor of minority groups, and more. Desai & Kroll (2017) continue this line and argue that transparency is insufficient regulation, offering other tools to test and evaluate algorithms. In contrast, Kleinberg, Ludwig, Mullainathan and Sunstein (2018) argue for a minimalistic approach to regulating algorithms, emphasizing the promise of algorithms in fighting discrimination given the right tools in place, and support transparency which they view as harmless relative to measures such as legally prohibiting the use of certain types of information as inputs. While there is no consensus as to what is the correct legal approach to regulating the use of algorithms, there is an implicit agreement that transparency is the least intrusive legal response, with the debate focused on whether stricter interventions are warranted.

Most of these (and other) studies, focus on the technological properties of algorithms, arriving at general results that are not sensitive to the widely varying economic functions that algorithms serve. Specifically, the legal scholarship is overwhelmingly focused on several specific types of algorithmic practices that are viewed as the most important, most notably algorithms that generate predictions based on race and their effect on racial equality. Notwithstanding the importance of fairness concerns, this paper abstracts from these important questions by assuming that consumers are homogenous in any aspect other than their willingness-to-pay for a product or service, which allows to focus on efficiency and consumer protection concerns overlooked by prior legal literature.[26] The analysis in the next section reveals, among other things, that transparency should not be taken lightly as the least intrusive legal response. The results of the model are that mandating disclosure of behavior-based pricing, given certain conditions, and even though it acts to eliminate consumer misperception in the market and help consumers make better decisions, could nonetheless result in harming consumers and decrease the overall welfare.

4. A Model of Behavior-Based Price Discrimination

4.1. Setup

The basic setup of the model is largely in the lines of Fundenberg and Villas-Boas (2006), that consider behavior-based pricing when consumers are strategic. The analysis is extended to account both for strategic consumers and myopic consumers in the same setting, in order to analyze the effect of algorithmic transparency. In each of two periods a single monopolistic seller of some good or service faces the same unit mass of potential consumers, where each consumer is uniquely identified as $i \in (0,1)$. The value that consumer i attaches to the product is $v_i \in [0, V]$. The unit mass of consumers are distributed along the range of valuations based on the density function $f(v)$ and the cumulative distribution function $F(v)$. v_i is time-invariant, meaning that each consumer has the same value for the product in every period.

The seller produces the product or service at a constant marginal cost that is assumed to be zero at every period, for simplicity. At the beginning of each period, the seller sets the price of the product, and records which consumers purchased the product and which did not purchase it. $a_{i,t} \in \{0,1\}$ indicates whether or not consumer i purchased the product in period t. The seller sets a single price in the first period $p_1 \in \mathbb{R}^+$, that is offered to all consumers. In every subsequent period, the seller can decide whether to set different prices to different segments of consumers, conditioned on their purchase history, where $p_2^{a_{i,1}} \in \mathbb{R}^+$ is the price of the product in the second period that is offered to consumer i, conditioned on her purchase history $a_{i,1}$. The seller cannot commit to pricing in advance – the strategies of all players must sustain a sub-game perfect equilibrium.

[26] Without losing sight of the fact that willingness-to-pay is often highly correlated with race, sex and other inherent features, which will weigh in when interpreting the results.

Each consumer either buys the product at the offered price or does not buy it. The consumer has no use for more than a single product in each period, and a product that is purchased in a certain period is consumed during that period, generating no value in future periods if not consumed. A consumer that purchases the product in period t receives a payoff of $v_i - p_t^{a_{i,t}}$ in that period, which assumes for simplicity that consumers have separable utility functions. Otherwise, they receive a payoff of zero, which can be interpreted either as not realizing the value from consuming the product or the normalized value of some outside option that the consumer has (for example, the net payoff from some substitute product, such as taking the bus instead of an Uber).

The product can be interpreted as any good or service, including (but not limited to) flight tickets, taxi rides, online groceries shopping, or personal loans, as long as the seller, provider of service or lender has access to some technology that lets it identify consumers and observe their transactional history $a_{i,t}$. In other words, the seller knows whether every potential customer (or "lead") purchased the product in the prior period. The framework does not apply, therefore, to brick-and-mortar sellers of goods that attach visible price labels to products, nor to sellers that cannot identify consumers at a reasonable effort. Observing the transactional history of consumers allows the seller to engage in price discrimination based on the transactional history of the consumer if it chooses to.

The seller and the consumers are risk-neutral and discount the value of future consumption by a rate of $\delta \in (0,1)$.

4.2. Analysis

The seller sets p_1 to all consumers in the first period, since at this point they are indistinguishable from one another in any relevant sense. In the second period, if consumers are identifiable, the seller can set two different prices, $p_t^{a_{i,1}} = p_2^1$ to consumers (if there are any) who purchased the product in the first period, and $p_t^{a_{i,1}} = p_2^0$ to consumers (if there are any) who did not purchase the product in the first period. Consumers in the second period can be described as either new or returning customers. While these are useful concepts, keep in mind that this is only one private case of behavior-based pricing. The two-period model is a simplified version of more realistic settings where sellers and consumers interact throughout multiple periods. When this is the case, two "returning customers" could nonetheless be offered different prices if their transactional histories exhibit different patterns.

Before proceeding to the main analysis, it is useful to present a fundamental feature of the decisions of consumers whether to purchase the product in the first period, as put forth in Lemma 1. In any Perfect Bayesian Equilibrium (PBE), the purchasing decision of consumers in the first period is monotonic in the valuation for the product. In other words, there exists a marginal consumer j, whose value for the product is denoted by $v_j \equiv \lambda$ such that every consumer whose value for the product is higher than λ will purchase it in the first period, and every consumer whose value for the product is lower than λ will not purchase it in the first period. This does not exclude the possibility that all consumers (or none) will purchase the product, but it means that *if* there is any separation of consumers in equilibrium – it will be characterized by a single valuation threshold.

Lemma 1. Monotonicity of Consumption Decisions

In any PBE, The purchasing decisions of consumers in period 1 are monotonic in their valuation of the product, defined as $\exists\lambda \in [0, V]: \forall i: \begin{cases} a_{i,1} = 0 & \text{if} \quad v_i < \lambda \\ a_{i,1} = 1 & \text{if} \quad v_i > \lambda \end{cases}$

Proof: In any PBE it must be that $\forall i: a_{i,1} = 1 \Leftrightarrow v_i - p_1 + \delta \max\{v_i - p_2^1, 0\} > \delta \max\{v_i - p_2^0, 0\}$ which is the constraint for individual rationality of purchasing the product in the first period. The definition of monotoncity formulated in lemma 1 is equivalent to the claim $\forall i \neq j: a_{i,1} = 1 \land a_{j,1} = 0 \Rightarrow v_j < v_i$. Suppose, for the sake of contradiction, that, w.l.o.g., $\exists i: a_{i,1} = 1 \land \exists j: a_{j,1} = 0 \land v_j \geq v_i$. $a_{i,1} = 1$ implies by IR constraint that $\max\{v_i + \delta v_i - p_1 - \delta p_2^1, v_i - p_1\} > \delta v_i - \delta p_2^0 \Leftrightarrow \max\{v_i - p_1 - p_2^1, v_i - \delta v_i - p_1\} > -\delta p_2^0$. $v_j > v_i$ implies that $\max\{v_j - p_1 - p_2^1, v_j - \delta v_j - p_1\} > \max\{v_i - p_1 - p_2^1, v_i - \delta v_i - p_1\}$ and by transitivity it follows that $\max\{v_j - p_1 - p_2^1, v_j - \delta v_j - p_1\} > -\delta p_2^0$ which implies $v_j - p_1 + \delta \max\{v_j - p_2^1, 0\} > \delta \max\{v_j - p_2^0, 0\}$, which by assumption 1 contradicts $a_{j,1} = 0$. □

The section will proceed as follows. First, it is useful to consider, as a benchmark, the static case where behavior-based pricing is either impossible, because the seller is unable to observe the transactional history of consumers, or effectively prohibited by law. Next, the seller will be able to observe the transactional history of consumers and set different prices based on it, but this ability to condition prices is entirely unknown to consumers, rendering them myopic, and behaving as if there is no possibility that future price will depend on their prior purchases. Lastly, the seller will be able to price based on prior behavior but it will be required to effectively disclose it, such that consumers are aware and informed of the possibility that the seller could choose to engage in behavior-based pricing. Consequently, consumers will be able to strategically decide whether to purchase the product based on the effect of their decision on future pricing. Ultimately, the desirability of banning behavior-based pricing or mandating the disclosure of behavior-based pricing will be explored from a welfare perspective.

4.2.1. Unidentifiable Consumers

It is useful to begin by considering the benchmark-setting where, for some reason, the seller cannot observe the prior purchases of customers or when the seller is legally obligated to set a single price to all consumers. This setting includes brick-and-mortar retail stores, payday loans,[27] or sellers that the law effectively restricts from price discriminating based on consumer behavior.

When the seller cannot condition prices on transactional history, it will set the price in each period as if there were no additional periods, based on its cost of production and the value of the product or service to consumers. Thus, in equilibrium, the monopolistic seller maximizes its profits without utilizing information on the consumers based on their prior behavior. In each period all consumers with $v_i > p_t = \lambda$ purchase the product. Knowing this, the seller sets the following identical price in each period $p_1^* = p_2^* = \arg\max_p p[1 - F(p)] \equiv p^S$, and in each period the consumers that purchase the product are those who value it above $\lambda^S = p^S$. Unless otherwise mentioned, $p[1 - F(p)]$ will be assumed to be concave, such that given any concrete distribution of consumers, p^S receives some single value. For example, for a uniform distribution, $p^S = \frac{V}{2}$. Hereinafter, the static monopolistic price p^S and the purchase decision λ^S will serve as a benchmark for the pricing strategy of the seller in the following sections that allow for behavior-based price discrimination.

4.2.2. Concealed Behavior-Based Pricing and Myopic Consumers

In this setting, the seller can set different prices in the second period conditioned on whether the consumer purchased the product in the first period. However, consumers are entirely unaware of the fact that their purchasing decisions could affect future pricing. Specifically, they are misled to believe that in every period the same price will be offered to everyone. This myopia could be the result of rationally bounded consumers that fail to anticipate or factor in the effect of their decisions on future pricing. Alternatively, and more relevant to the motivation of this paper, it could represent a successful attempt on part of the seller to mislead consumers to believe that it does not possess the technology required to engage in behavior-based pricing. This misleading could be a result of proactive efforts taken by the seller, or simply the fact that consumer awareness of ways in which businesses use algorithms to set prices tends to lag behind the technological state-of-the-art.

[27] Information on payday loans is generally not reported to the large credit scoring agencies, unless the debt is sold to a collection agency. Therefore, payday loans typically affect credit score (and through it the terms of future loans) less then standard loans.

The seller may now set different prices in the second period to the two segments of consumers. Returning customers are offered p_2^1 and new customers are offered p_2^0. The seller knows whether each customer purchased the product in the first period, and this is captured by the value λ that separates the consumers into the two segments. Because the second period is the last, consumers will purchase the product in the second period whenever their value from it is higher than the price offered to them, as given by:

$$(1) \qquad \forall i: a_{i,2} = 1 \Leftrightarrow v_i - p_2^{a_{i,1}} > 0$$

Anticipating the decisions of consumers, the seller will set each of the two second-period prices to maximize its profits from serving the segment of consumers to whom that price is offered. Concretely, each of the two optimal second-period prices can be formulated as a function of λ, as follows:

$$(2) \qquad p_2^{0*}(\lambda): \arg\max_{p_2^0} \pi_2^0 = \arg\max_{p_2^0} \left[p_2^0 \left(F(\lambda) - F(p_2^0)\right)\right]$$

$$(3) \qquad p_2^{1*}(\lambda): \arg\max_{p_2^1} \pi_2^1 = \arg\max_{p_2^1} \left[p_2^1 \left(1 - F(\max\{p_2^1, \lambda\})\right)\right] = \max\{\lambda, p^S\}$$

Equation (2) defines p_2^{0*} to be the price that maximizes the profits from serving the segment of consumers that did not purchase the product in the first period – those who value the product below λ. Equation (3) defines p_2^{1*} to be the price that maximizes the profits from serving the segment of returning customers. The consumers that purchase the product in the second period at a price p_2^{1*} are those who value the product above λ (otherwise, they are not offered this price) *and* above p_2^{1*} (otherwise, they are better off not purchasing it). This somewhat unique feature allows deriving an explicit function for p_2^{1*}, which is simply the greater of λ or p^S.[28]

In the first period, consumers decide whether to purchase the product based solely on maximizing their payoff in that single period, i.e., they purchase it whenever $v_i \geq p_1$. This represents the assumption that consumers are unaware of the fact that the seller may set future prices based on purchase history. Therefore, consumers do not take into account that the decision to purchase may affect the price that they will be offered in the future. Formally, the separation of consumers into two segments, as induced by the price that the seller sets in the first period, is simply given by:

---

[28] To see this, recall that the best that the seller can do vis-à-vis a single segment of consumers, by definition, if it can, is to set a price of p^S. However, this is feasible only if the segment is large enough to include all those who value the product above p^S (i.e., if $\lambda \leq p^S$). If, on the other hand, the segment is smaller, such that $\lambda > p^S$, then the best that the seller can do is set the price at λ and serve all consumers in the segment. Formally, it must be that (1) $p_2^{1*} \geq \lambda$, otherwise increasing p_2^1 will strictly increase profits. (2) Because $p[1 - F(p)]$ is maximized at p^S, $p_2^{1*} = p^S$, unless condition (1) does not hold, which happens if $p^S < \lambda$. In that case, because $p[1 - F(p)]$ is concave, $p_2^{1*} = \lambda$. Therefore, p_2^{1*}=$\max\{\lambda, p^S\}$.

$$(4) \qquad \lambda(p_1) = p_1$$

The seller anticipates that the price in the first period will not only affect its profits in the first period but also generate information on the valuation of consumers that can be used in the future to set different prices. Taking this into account, the seller sets the first-period price by maximizing the following profit function (where λ, p_2^{0*} and p_2^{1*} are all functions of p_1, either directly or indirectly):

$$(5) \qquad \max_{p_1} \pi_1 + \delta(\pi_2^{0*} + \pi_2^{1*}) = p_1(1 - F(\lambda)) + \delta p_2^{0*}(F(\lambda) - F(p_2^{0*})) + \delta p_2^{1*}(1 - F(p_2^{1*}))$$

The first term is the seller's profits in the first period, when the product is purchased by all consumers with $v_i \geq \lambda$. The second term is the profits in the second period generated from serving the segment of consumers that did not purchase the product in the first period. The quantity purchased by new customers $(F(\lambda) - F(p_2^{0*}))$ increases in λ, and therefore in p_1, since λ determines the upper bound, in term of valuation, of the segment of consumers that did not purchase the product in the first period and are now offered a price of p_2^0. The third term is the profits in the second period generated from serving the segment of consumers that purchased the product in the first period. Recall that $p_2^{1*} = \max\{\lambda, p^S\}$, such that an initial result is that in any PBE, it must be that $p_2^{1*} \geq p_1^* = \lambda$. The main results are presented in Proposition 1 and illustrated in Figure 1.

Proposition 1. Concealed Behavior-Based Pricing (with myopic consumers)

There is a unique PBE where the monopolist seller will set prices and consumers will purchase the product as characterized by the following:

1. The first-period price is $p_1^* = p^M > p^S$ and all consumers with $v_i > \lambda^M = p^M > \lambda^S$ purchase the product.
2. The second-period price for consumers that purchased in the first period is $p_2^{1*} = p^M > p^S$.
3. The second-period price for consumers that did not purchase in the first period is $p_2^{0*} < p^S < p^M$.

Proof: Denoting $\lambda' \equiv \frac{\partial \lambda}{\partial p_1}$ and $p_2^{0*\prime} \equiv \frac{\partial p_2^{0*}}{\partial p_1}$, and assuming (to be sustained in equilibrium) that $p_2^{1*} = \max\{\lambda, p^S\} = \lambda$, the F.O.C. of the maximization problem formulated in equation (5) is given by the following condition: $\frac{\partial \pi}{\partial p_1} = 1 - F(\lambda) - p_1^* f(\lambda)\lambda' + \delta\big[p_2^{0*\prime} F(\lambda) + p_2^{0*} f(\lambda)\lambda' - p_2^{0*\prime} F(p_2^{0*}) - p_2^{0*} f(p_2^{0*}) p_2^{0*\prime} + \lambda' - \lambda' F(\lambda) - \lambda f(\lambda)\lambda'\big] = 0$. Notice that from equation (4) $\lambda(p_1) = p_1$ and $\lambda' = \frac{\partial \lambda}{\partial p_1} = 1$. Furthermore, from equation (2) it must hold in PBE that $F(\lambda) - F(p_2^{0*}) - p_2^{0*} f(p_2^{0*}) = 0$. Plugging these into the F.O.C. reduces it to: $1 - F(\lambda) - \lambda f(\lambda) + \delta[p_2^{0*} f(\lambda) + 1 - F(\lambda) - \lambda f(\lambda)] = 0$, such that $\lambda^* = p_1^* = p^M$ solve this equation. Suppose, for the sake of contradiction, that $p_1^* = p^M = p^S$. By definition of p^S, $1 - F(p_1^*) - p_1^* f(p_1^*) = 0$ if $p_1^* = p^S$. Furthermore, since $p[1 - F(p)]$ is concave, $1 - F(p_1^*) - (p_1^* - p_2^{0*}) f(p_1^*) > 0$ if $p_1^* = p^S$, such that $\frac{\partial \pi}{\partial p_1}(p_1 = p^S) > 0$. Therefore, it must be that $p_1^* = p^M > p^S$. Since $\lambda^* = p^M > p^S$, then by equation (3) it is sustained in equilibrium that $p_2^{1*} = \lambda^* = p^M$. Finally, recall that from equation (2) we get that $\frac{\partial \pi_2^0}{\partial p_2^0} = F(\lambda) - F(p_2^0) - p_2^0 f(p_2^0)$. Since p^S is defined as

$1 - F(p^S) - p^S f(p^S) = 0$, then it must be that for $p_2^0 = p^S$, $\frac{\partial \pi_2^0}{\partial p_2^0} < 0$ for any $\lambda < 1$. Therefore, in equilibrium, it must be that $p_2^{0*} < p^S$. □

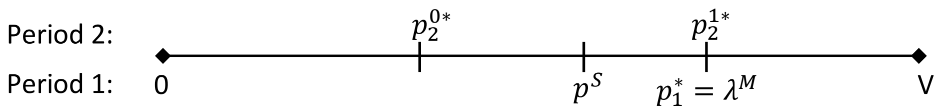

Figure 1: Concealed Behavior-Based Pricing (with myopic consumers)

Intuitively, increasing p_1 from the naïvely profitable level of p^S will decrease the seller's profits (and the consumer surplus) in the first period. However, it will increase λ^S, which means that the segment of consumers that do not purchase the product in the first period, and are therefore offered p_2^{0*} in the second period, will be larger and have a higher average willingness-to-pay. When confronted with a segment of consumers with higher valuation for the product, the seller can set p_2^{0*} in a way that generates more profits from serving that segments. In other words, increasing p_1 will allow the seller to use the information generated by the market segmentation that it induces, λ^S, to profitably increase the price offered to the segment of new customers. Since $\lambda^S > p^S$, it follows that $p_2^{1*} > p^S$, since it does not make sense to set a price that is lower than the willingness-to-pay of the consumer that values the product the least in that segment. The result is that behavior-based pricing, when consumers are myopic, leads to a price increase both in the first period and in the second period to consumers who previously purchased the product. Notice that the informational benefit, allowing the seller to set a more profitable p_2^{0*} to new customers, comes at a cost. Since p^S is the single-period profit maximizing price, and departure of both p_1^* and p_2^{1*} further away from p^S, leads to a decrease in profits gained vis-à-vis each of these segments of consumers.

In the example of a uniform distribution (where $\forall p: f(p) = \frac{1}{V}, F(p) = \frac{p}{V}$), it is easy to see that the seller sets a price in the first period of $p_1^* = \lambda^M = p^M = \frac{2V(1+\delta)}{4+3\delta}$ (which, for any δ, is higher than the single-period price of $p^S = \frac{V}{2}$). In the second period, the seller sets $p_2^{1*} = p^M = \frac{2V(1+\delta)}{4+3\delta}$ to returning customers and $p_2^{0*} = \frac{V(1+\delta)}{4+3\delta}$ to new customers (which, for any δ, is lower than the singe-period price of $p^S = \frac{V}{2}$)).

To summarize, when the seller can engage in behavior-based pricing without the consumers being aware of it, the seller will find it profitable to do so by setting a higher price in the first period (relative to the monopoly price in the single-period static benchmark). Furthermore, in the second period it will set the same higher price for returning customers, but a lower price for new customers. This result already reveals a deviation from the second-best efficient outcome of static monopolistic pricing, since even fewer consumers will be served in the first period, due to the higher price. However, more consumers will be served in the second period due to the differentiated pricing, as predicted by the classical accounts of price discrimination.

4.2.3. Behavior-Based Pricing with Informed Consumers

When consumers are made informed of the pricing strategy of the seller, and if they are strategic, then the seller cannot simply assume that a consumer will purchase the product in the first period whenever they value it more than its price, such that $\lambda = p_1$, as in the previous settings. Some consumers with valuation higher than p_1 may be better off holding back from an otherwise efficient purchase to secure a lower price in the future, inducing a market segmentation with a higher valuation threshold $\lambda > p_1$.

Following Fudenberg and Villas-Boas (2006), deriving the behavior of the seller and consumers in a PBE in this setting is done by backward induction. In the second (and last) period, when there are no future payoffs to be considered, consumers will purchase the product as before, whenever the price offered to them is lower than the value of the product, as given by equation (1). The seller will set the prices in the second period separately for each segment of consumers, those who previously purchased the product and those who did not, by solving two single-period maximization problems for each segment of consumers separately. As before, given a market segmentation characterized by λ, the seller will set the second-period prices p_2^0 and p_2^1 as given by equations (2) and (3). The meaningful difference from the previous settings, therefore, is not the behavior of the parties in the second period, but rather in how they will make their decisions in the first period to induce a segmentation characterized by λ.

Informed and strategic consumers will purchase the product in the first period if the first-period payoff from purchasing the product combined with their expected payoff from being offered the product in the second period at a price p_2^1 is greater than their expected payoff from being offered the product in the second period at a price p_2^0. This condition in a two-periods setting is given by:

$$(6) \qquad \forall i: a_{i,1} = 1 \Leftrightarrow v_i - p_1 + \delta \max\{v_i - p_2^{1*}, 0\} > \delta \max\{v_i - p_2^{0*}, 0\}$$

Recall that the expected price for returning customers is $p_2^{1*}(\lambda) = \max\{\lambda, p^S\}$. Therefore, it is clear that the marginal consumer, whose value for the product is $v_i = \lambda$, will not expect to obtain a positive payoff from purchasing the product in the second period if she decides to purchase the product in the first period, since $\lambda - p_2^{1*} \leq 0$, which implies that for the marginal consumer $\delta \max\{\lambda - p_2^{1*}, 0\} = 0$. Furthermore, it is clear that $p_2^{0*} < \lambda$, since $p_2^{0*} = \lambda$ would result in zero profits in the second period from (not) serving the segment of consumers that did not previously purchase the product. Therefore, the marginal consumer will obtain a positive payoff of $\lambda - p_2^{0*}$ in the second period if she does not purchase the product in the first period, implying that $\delta\max\{\lambda - p_2^{0*}, 0\} = \delta(\lambda - p_2^{0*})$. Therefore, applying equation (6) for the marginal consumer allows to derive the specific function for λ as a function of p_1:

$$(7) \qquad \lambda(p_1) = \frac{p_1 - \delta p_2^{0*}}{1-\delta}$$

Evidently, the segmentation threshold increases in the first-period price. Moreover, it decreases in the expected second-period price to consumers that do not purchase the product in the first period. Intuitively, if avoiding the purchase is expected to lead to a lower price in the future, then more people will be better off holding off the otherwise efficient purchase. Notice that equation (7) is sufficient to determine that behavior-based pricing will occur in equilibrium, in the sense that $p_1^* > p_2^{0*}$. If we would assume that in equilibrium the strategy of the seller is to set a price in the first period that is equal to the price offered to new customers in the second period, $p_1^* = p_2^{0*} = a$, then by equation (7) it would induce a segmentation of $\lambda = a$. The reason is that consumers would have no reason to hold off an efficient purchase in expectation of a price cut. However, this pricing strategy is not credible. Clearly, once the seller gets to subgame that begins at the second period, when $\lambda = a$, it will find it profitable to deviate from its strategy and set $p_2^{0*} < p_1^* = a$. The first-period pricing decision that is part of a strategy in a subgame perfect Bayesean equilibrium is determined by the following maximization problem:

$$(8) \qquad \max_{p_1} \pi_1 + \delta(\pi_2^{0*} + \pi_2^{1*}) = p_1(1 - F(\lambda)) + \delta p_2^{0*}(F(\lambda) - F(p_2^{0*})) + \delta p_2^{1*}(1 - F(p_2^{1*}))$$

For ease of exposition, this maximization problem is formulated identically to the one given by equation (5), when the the possibility of behavior-based pricing was concealed from consumers. Crucially, however, when the seller sets p_1 it will now have to consider what market segmentation λ the price will induce. The market segmentation is not only a function of p_1 but also it is also determined by the expected future price to new customers p_2^{0*}, which in itself is simultaneously a function of λ that is sustained in equilibrium. The main results are presented in Proposition 2.

Proposition 2. Behavior-Based Pricing (with informed and strategic consumers)
There is a unique PBE where the monopolist seller will set prices and consumers will purchase the product as characterized by the following:

1. The first-period price is $p_1^* = p^B$ and all consumers with $v_i > \lambda^B > \lambda^S$ purchase the product.
2. The second-period price for consumers that purchased in the first period is $p_2^{1*} = \lambda^B > p^B$.
3. The second-period price for consumers that did not purchased in the first period is $p_2^{0*} < p^B, p^S$.

Proof: from proposition 1 we know that $\frac{\partial \pi}{\partial p_1} = 1 - F(\lambda) - p_1 f(\lambda)\lambda' + \delta\lambda'[p_2^{0*} f(\lambda) + 1 - F(\lambda) - \lambda f(\lambda)] = 0$. Equation (7) can be rearranged as $\lambda(1-\delta) = p_1 - \delta p_2^{0*}$, such that the F.O.C. can be rewritten as: $1 - F(\lambda) - f(\lambda)\lambda(1-\delta)\lambda' + \delta\lambda'[1 - F(\lambda) - f(\lambda)\lambda] = 0$. Using the generalized chain rule, it is possible to differentiate the implicit function for λ given in equation (7). Specifically, $\lambda' = \frac{\partial \lambda}{\partial p_1} = \frac{1}{1-\delta+\delta\frac{\partial p_2^{0*}}{\partial \lambda}}$, which is a function of λ itself. Therefore, $\lambda^* = \lambda^B$ is defined as solving the F.O.C., and p_2^{0*} is given directly by equation (2). Rearranging equation (7) then gives $p_1^* = p^B = (1-\delta)\lambda^B + \delta p_2^{0*}$. Suppose, for the sake of contradiction, that $\lambda^B = \lambda^S = p^S$. By definition of p^S, $1 - F(\lambda^B) - \lambda^B f(\lambda^B) = 0$. Next, it is easy to verify by implicit differentiation that if $p[1-F(p)]$is concave then $\frac{\partial p_2^{0*}}{\partial \lambda} > 0$, and therefore $\lambda'(1-\delta) = \frac{1-\delta}{1-\delta+\delta\frac{\partial p_2^{0*}}{\partial \lambda}} \in (0,1)$. Therefore, and by definition of p^S, $1 - F(\lambda^B) - f(\lambda^B)\lambda^B(1-\delta)\lambda' > 0$. The result is that if $\lambda^* = \lambda^B = \lambda^S = p^S$ then $\frac{\partial \pi}{\partial p_1} > 0$, and therefore it must be that in equilibrium $\lambda^B > \lambda^S = p^S$. Since $\lambda^* = \lambda^B > p^S$, then by equation (3) it is sustained in equilibrium that $p_2^{1*} = \lambda^B$. Finally, as shown above, in a PBE it must be that $p_2^{0*} < p_1^* = p^B$, and it is easy to see from equation (7) that $p^B < \lambda^B = p_2^{1*}$. Finally, recall that from equation (2) we get that $\frac{\partial \pi_2^0}{\partial p_2^0} = F(\lambda) - F(p_2^0) - p_2^0 f(p_2^0)$. Since p^S is defined as $1 - F(p^S) - p^S f(p^S) = 0$, then it must be that for $p_2^0 = p^S$, $\frac{\partial \pi_2^0}{\partial p_2^0} < 0$ for any $\lambda < 1$. Therefore, in equilibrium, it must be that $p_2^{0*} < p^S$. □

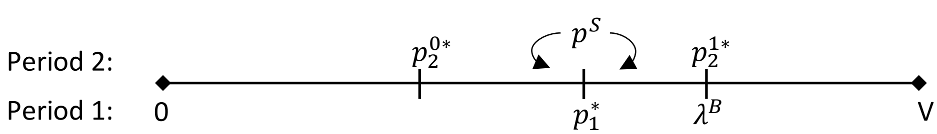

Figure 2: Behavior-Based Pricing (with informed and strategic consumers)

The results summarized in proposition 2 are consistent with Fudenberg and Villas-Boas (2006) that considered a similar setting. As in the case where behavior-based pricing was concealed, the seller is driven to increase the first-period price p_1 above its naïvely profitable level of p^S, to induce a market segmentation that will allow it to gain more profits in the second period. However, crucially, this is only part of the story. Now that consumers are aware of this pricing scheme, some of them will strategically refrain from purchasing in the first period even if they value the product more than its price, in order to secure a lower price in the future. The strategic response by some of the consumers acts to increase the market segmentation threshold λ upwards, everything else equal. Consequently, however, there is now a second force that drives the seller to the opposite direction – to decrease the first-period price. The reason is that now the seller has to offer a better deal for consumers in order to appeal to some of the strategic ones and induce a market segmentation that is closer to the desired level. In other words, when consumers behave strategically, the decision to purchase the product generates an even stronger signal as to the willingness-to-pay of those that purchase it, because the consumers who purchase have to value the product sufficiently higher than its price, such that purchasing pays off despite the future penalty in the form of a higher price. The seller can, therefore, more safely deduce that consumers that purchase the product in the first period (despite being aware on its effect on future prices) value it more relative to when the pricing scheme was concealed or forbidden. The outcome is that it is inconclusive whether the seller will set a lower or higher price in the first period. However, and more importantly from an efficiency perspective, it is certain that whether the price will be lower or higher, it will be set such that fewer consumers will purchase the product, $\lambda^B > \lambda^S$.

As a side note, and as previously noted in the literature, if the seller was able to commit to future pricing, it would be better off by committing not to price discriminate, if the alternative is to disclose behavior-based pricing. Specifically, the seller would set the same price of p^S in each period and to all consumers, disregarding past purchases. Somewhat counterintuitively, however, the fact that the seller would prefer not to price discriminate if the alternative is to disclose does not mean that it is not efficient, as many consumers benefit from the lower prices induced by price discrimination. Moreover, it does not say anything as to whether disclosure is desirable, given that behavior-based pricing is practiced, which will be explored in the next section. In any case, the plausibility of adopting commitment mechanisms is not explored in this paper, which is focuses on disclosure rather than voluntary renouncement. However, if and once disclosure is mandated, it would undoubtedly be pertinent to further our understanding of whether and how sellers may be able to profitably commit to pricing in a world where algorithms often operate like a "black box," which is not fully transparent even to its engineers.

In the example of a uniform distribution (where $\forall p\colon f(p) = \frac{1}{V}, F(p) = \frac{p}{V}$), it is easy to verify that the seller sets a price in the first period of $p_1^* = p^B = \frac{V(4-\delta^2)}{2(4+\delta)}$. This price is lower (for any δ) than the single-period price of $p^S = \frac{V}{2}$. In the example of a uniform distribution, consumers are relatively dispersed along the dimension of the value for the product. Intuitively, this makes it more profitable for the seller to try and appeal to a larger number of consumers by reducing both the first-period price and, consequently, the second-period price offered to new customers (rather than trying to extract more value from a narrower subset of consumers). Despite the lower price, it induces a segmentation threshold of $\lambda^B = \frac{V(2+\delta)}{4+\delta}$ that is higher (for any δ) than $\lambda^S = \frac{V}{2}$.

To summarize, when consumers are made aware of the fact that the seller is engaging in behavior-based pricing, some of them who value the product more than its price avoid an otherwise efficient purchase in order to secure an even lower price in the future. Anticipating this, the seller may decide to either decrease or increase the first-period price, but it will never decide to decrease it to the extent that would induce the same volume of purchases as would have been made if pricing was not behavior-based, making it potentially inefficient. The next section will combine the accumulated insights to explore whether mandating effective disclosure of behavior-based pricing is desirable from a welfare perspective.

4.2.4. Welfare Analysis of Behavior-Based Pricing and Mandated Disclosure

In each of the previous two sections, behavior-based pricing – whether concealed or not – was characterized relative to the benchmark-setting where behavior-based pricing is impossible or forbidden. A natural question that arises is whether behavior-based pricing should be permitted (and even encouraged) or rather deemed as illegal price discrimination. On the one hand, with behavior-based pricing, fewer consumers purchase the product in the first period, $\lambda^M, \lambda^B > \lambda^S$, resulting in reduced efficiency. On the other hand, in the second period, all consumers who value the product above p_2^0 purchase the product (whether they buy it at p_2^0 or p_2^1 does not currently matter). As we saw earlier, with behavior-based pricing $p_2^0 < p^S$, such that more consumers purchase the product in the second period.

The result that more consumers purchase the product in the second period is the standard result when sellers can engage in price discrimination. However, in the case of behavior-based pricing, the ability to efficiently price discriminate comes at a cost – serving fewer consumers in the first period. In other words, the information that the seller uses to price discriminate is not exogenous, as most accounts of price discrimination assume, but rather it is generated endogenously at a cost. The trade-off between the costs of obtaining information on consumers (by forgoing efficient sales in the first period) and the value from serving more consumers in the second period is inconclusive, and should be examined based on the characteristics of any specific case. Generally, the welfare in any given one of the three legal regimes analyzed $L \in \{S, M, B\}$ – behavior-based pricing is banned, permitted and concealed, or permitted and disclosed – can be formulated as $W^L = \delta \int_{p_2^{0L}}^{V} vf(v)dv + \int_{\lambda^L}^{V} vf(v)dv$, which is simly the aggregated willingness-to-pay for all consumers that purchase the product in both periods. Using this expression, the efficiency gains (or losses) from permitting behavior-based pricing (either with or without disclosure) can be formulated as $\Delta W^{M/B} = \delta \int_{p_2^{0M/B}}^{p^S} vf(v)dv - \int_{p^S}^{\lambda^{M/B}} vf(v)dv$. This formulation does not offer further insights beyond the intuition described above, that the welfare effect depends on whether sufficiently more consumers purchase the product in the second period relative to the fewer that purchase it in the first period. However, it can be used for empirically testing the efficiency of behavior-based pricing given the existence of the relevant pricing data (and plausible assumptions as to the extent to which the valuation for the product varies among consumers).

In the example of a uniform distribution (where $\forall p: f(p) = \frac{1}{V}, F(p) = \frac{p}{V}$), and normalizing, for simplicity of presentation, $V = 1$, the overall welfare when consumers are unidentifiable is $W^S = \frac{3}{8}(1 + \delta)$. Furthermore, the overall welfare when behavior-based pricing is permitted and consumers are either myopic or informed is, respectively, $W^M = \frac{1}{2}(1 + \delta) - \frac{(1+\delta)^2}{2(4+3\delta)^2}(4 + \delta)$ and $W^B = \frac{1}{2}(1 + \delta) - \frac{(2+\delta)^2}{8(4+\delta)}$. It is easy to verify that $W^M > W^S$ and that $W^B > W^S$ for any $\delta \in (0,1)$, such that banning behavior-based pricing is inefficient, whether it is concealed from consumers or not (it is also the case that $W^M > W^B$, but this will be explored separately in the remainder of the section).

Surely, in many cases where consumers are not uniformally distributed, it may be that behavior-based pricing is inefficient, or that its efficiency depends on δ. Notwithstanding the importance of conducting a welfare analysis that is sensitive to the particularities of any market, entirely forbidding the use of behavior-based pricing seems somewhat implausible. Behavior-based pricing does not use directly protected characteristics such as race or sex, and in the algorithmic age it may not be fully transparent even to the businesses that use it. Moreover, in practice, where sellers and consumers often interact multiple times, it is possible that over time the efficiency gains from customized pricing will overshadow the initial costs of generating the information required for it. Therefore, banning or permitting behavior-based pricing seems overly simplistic relative to the wide range of possible legal responses to it.

Therefore, the remainder of the section will focus on exploring an intermediate but common legal intervention – whether sellers that engage in behavior-based pricing should be required to make this feature of their pricing scheme public. In other words, whether it is desirable to mandate disclosure of behavior-based pricing practices or the transparency of algorithms that fulfill this function. The measure of welfare considered here is the overall welfare of consumers and the seller combined. The question of whether rulemaking should be designed to increase the overall welfare or the welfare of consumers is a heated debate. However, as will be shown, the interpretation of the results are not sensitive to this question, because in the context of disclosing behavior-based pricing, consumer surplus and overall welfare are mostly aligned.

First, it is useful to observe that mandated disclosure will necessarily harm the seller by decreasing its profits. This finding is consistent with the intuition that that sellers should be expected to conceal behavior-based pricing once they engage in it, and cannot be relied on to voluntary disclose behavior-based pricing (unlike surge pricing, for example, which performs better for the benefit of sellers in balancing supply and demand when consumers are aware and responsive to it). To see this, notice that when consumers are myopic, the seller could, in principle, set $p_1^M = \lambda^B > p_1^B$. If the seller does that, the same number of consumers will purchase the product in the first period as when consumers are strategic and the seller maximizes its profits, but the first-period price is higher. Therefore, the first-period profits when consumers are myopic and the price is set to be $p_1^M = \lambda^B$ are strictly higher than the maximal first-period profits when consumers are strategic. Furthermore, notice that when consumers are myopic, $\lambda^M = p_1^M$, which means that $\lambda^B = \lambda^M$. Therefore, the second-period prices are identical in this case. The conclusion is that when consumers are myopic and $p_1^M = \lambda^B$, the profits that the seller make are strictly higher than the maximal profits when consumers are strategic. Therefore, by the weak axiom of revealed preferences, if the seller sets any price other than $p_1^M = \lambda^B$ it will further increase its profits. The conclusion is that mandated disclosure will certainly decrease the profits of the seller.

As mentioned earlier, it has already been found in prior literature that the seller will rather commit not to price discriminate, if it was possible, than to engage in behavior-based pricing when consumers are strategic. Moreover, behavior-based pricing when consumers are myopic is trivially better for the seller than when it is prohibited, because the seller could always set all prices to be p^S and obtain the same profits, but it is more profitable to price discriminate. Combining these two observations together with the finding in the previous paragraph results in an interesting and counterintuitive outcome. Denoting the optimal profits of the seller given one of the three legal regimes analyzed $L \in \{S, M, B\}$ as π^L, we get that $\pi^M > \pi^S > \pi^B$. This means that a seller prefers that behavior-based pricing would be legal, but only if it can mislead consumers into believing that it is not practiced ($\pi^M > \pi^S$). However, if a seller is obligated to disclose engaging in behavior-based pricing, than it is better off if it would be made illegal altogether ($\pi^S > \pi^B$), because the legal prohibition is assumed to work as an effective commitment mechanism.

An interim conclusion is that sellers are expected to object to mandated disclosure, ruling out the possibility of voluntary transparency, but this does not determine yet whether disclosure is desirable from a welfare perspective. As noted earlier, given some market segmentation characterized by λ, the seller sets the same prices in the second period, whether behavior-based pricing is disclosed or not, as given by equations (2) and (3). The consumers that purchase the product in the second period are all those who value it above p_2^0, which is therefore an indicator of the overall welfare in the second period.[29] By implicitly differentiating the condition for the optimal p_2^{0*} as given by equation (2), we get that the sensitivity of p_2^{0*} to changes in λ is given by $\frac{\partial p_2^{0*}}{\partial \lambda} = \frac{f(\lambda)}{2f(p_2^{0*}) + p_2^{0*} f\prime(p_2^{0*})}$. By the assumption that $p[1 - F(p)]$is concave it is easy to verify that $\frac{\partial p_2^{0*}}{\partial \lambda} > 0$ for any λ. Therefore, a higher segmentation threshold λ leads to a higher p_2^{0*}, which results in fewer consumers purchasing the product and lower welfare in the second period. Moreover, in the first period, a higher λ trivially leads to lower welfare, since by definition, it means that fewer consumers purchase the product.

Therefore, when the seller engages in behavior-based pricing, the overall two-period welfare is a function that is decreasing in λ. The implication for a welfare analysis is that while the segmentation of the market λ is not a quantitative measure for welfare, it can conveniently be used to compare which of any two states corresponds with higher overall welfare. Denoting the overall welfare when behavior-based pricing is concealed as W^M and the overall welfare when consumers are informed of the behavior-based pricing as W^B, Corollary 1 presents the necessary and sufficient condition under which either of the states is superior over the other.

Corollary 1. Welfare with Behavior-Based Pricing (concealed versus disclosed)

$$W^B < W^M \Leftrightarrow \lambda^B > \lambda^M \Leftrightarrow \frac{\partial p_2^{0*}}{\partial \lambda}(\lambda = \lambda^M) > \tau$$

Where:

$$\tau = -\frac{1}{\delta} \times \frac{1 - F(\lambda^M) - f(\lambda^M)\lambda^M}{1 - F(\lambda^M)}$$

Proof: from proposition 2 we know that $\frac{\partial \pi^B}{\partial p_1} = 1 - F(\lambda) - p_1 f(\lambda)\lambda' + \delta\lambda'[p_2^{0*} f(\lambda) + 1 - F(\lambda) - \lambda f(\lambda)] = 1 - F(\lambda) + \lambda'[\delta(1 - F(\lambda)) - f(\lambda)(\delta\lambda - \delta p_2^{0*} + p_1)]$. Rearranging equation (7) gives $p_1 = (1 - \delta)\lambda^B + \delta p_2^{0*}$, and implicitly differentiating λ w.r.t. p_1 gives $\lambda' = \frac{\partial \lambda}{\partial p_1} = \frac{1}{1 - \delta + \delta \frac{\partial p_2^{0*}}{\partial \lambda}}$. By plugging these we get $\frac{\partial \pi^B}{\partial p_1} = 1 - F(\lambda) + \frac{1}{1 - \delta + \delta \frac{\partial p_2^{0*}}{\partial \lambda}}[\delta(1 - F(\lambda)) - f(\lambda)\lambda]$. Therefore, $\frac{\partial \pi^B}{\partial p_1} > 0$ when $(1 - \delta + \delta \frac{\partial p_2^{0*}}{\partial \lambda})(1 - F(\lambda)) + [\delta(1 - F(\lambda)) - f(\lambda)\lambda] > 0$, which is equivalent to, after rearranging $\frac{\partial p_2^{0*}}{\partial \lambda} > \frac{-(1 - F(\lambda) - f(\lambda)\lambda)}{\delta(1 - F(\lambda))}$. Therefore, if $\lambda^B = \lambda^M$ then $\frac{\partial \pi^B}{\partial p_1}(\lambda = \lambda^M) > 0 \Leftrightarrow \frac{\partial p_2^{0*}}{\partial \lambda}(\lambda = \lambda^M) > -\frac{1}{\delta} \times$

[29] Recall that in any PBE $p_2^{1*} = \lambda$, so there is no "gap" of consumers that are not served in the second period. All the consumers who value the product above λ are offered p_2^{1*} and purchase the product, and all consumers that value the product below λ are offered p_2^{0*}, and those who value the product above p_2^{0*} purchase it. Therefore, p_2^{0*} is a sufficient indicator for the overall number of consumers that purchase the product in the second period, from both segments. Since welfare is determined only by the number of people that purchase the product (and their value from it), and not the prices that they pay (i.e., p_2^{1*} or p_2^{0*}), then p_2^{0*} is a sufficient indicator to the welfare in the second period.

$\frac{1-F(\lambda^M)-f(\lambda^M)\lambda^M}{1-F(\lambda^M)} \equiv \tau$. This means that the optimal p_1^B is higher than the price that induces $\lambda = \lambda^M$. Since $\frac{\partial \lambda}{\partial p_1} > 0$, it must be that the optimal segmentation $\lambda^B > \lambda^M$. As shown above, this is a necessary and sufficient condition for $W^B < W^M$. □

As formulated in Corollary 1 (although, there are other possible formulations), the welfare effect of mandating the disclosure of behavior-based pricing can be understood through the lens of how sensitive is the second-period price offered to new customers p_2^{0*} to changes in the market segmentation λ. Intuitively, when the distribution of consumers is such that a larger segment of new customers in the second period (higher λ) leads to a sharper increase of the price offered to them, p_2^{0*}, then, in order to receive these potential profits, everything else equal, the seller will be willing to increase λ by increasing the first-period price p_1. Increasing p_1 and inducing a higher λ decreases the profits both in the first period and in the second period vis-à-vis returning customers. However, when $\frac{\partial p_2^{0*}}{\partial \lambda}$ is high enough, above τ, the seller will be made better off by forgoing these profits in exchange for higher profits received from serving the segment of new customers.

This result highlights a counterintuitive feature of behavior-based pricing. Requiring sellers to reveal the fact that they use algorithms to price discriminate based on past behavior may (though not necessarily) lead to lower welfare. This effect of disclosure is not *despite* the fact that consumers are being informed, but precisely because they are. Informed consumers may choose to hold off from otherwise efficient purchases, in anticipation of future lower prices. Anticipating this, the seller is very likely to decrease the price in the first period, to better appeal to consumers. However, this effect may be insufficient to offset the inefficiencies caused by the strategic purchasing decisions of consumers.

The threshold τ, above which $\frac{\partial p_2^{0*}}{\partial \lambda}$ dictates that disclosure is welfare decreasing, does not take imaginary high values, such that this possibility is rendered merely theoretic. For example, in the case of uniform distribution, $\frac{\partial p_2^{0*}}{\partial \lambda} = \frac{1}{2}$ for any level of λ. The sensitivity threshold is $\tau = \frac{1}{2+\delta}$, which is lower than $\frac{1}{2}$ for any δ. Therefore, for a uniform distribution, $\lambda^B > \lambda^M$ and therefore $W^B < W^M$, meaning that disclosure will decrease welfare. The same conclusion can be reached equivalently by explicitly finding the welfare under each legal regime. As was shown earlier, with a uniform distribution $W^M = \frac{1}{2}(1+\delta) - \frac{(1+\delta)^2}{2(4+3\delta)^2}(4+\delta)$ and $W^B = \frac{1}{2}(1+\delta) - \frac{(2+\delta)^2}{8(4+\delta)}$, and it is easy to verify that $W^B < W^M$ for any $\delta \in (0,1)$. With other types of distributions, the conclusion will depend on factors such as how many consumers are there that value the product in the relative vicinity of the prices set by the seller, as well as the discount factor δ. The model can be used to empirically predict the desirability of disclosure mandates given data on the characteristics of the market.

Finally, as mentioned above, consumer surplus is mostly aligned with the overall welfare effect of disclosure. Therefore, the suggested indicator for whether a disclosure mandate is desirable is not significantly sensitive to the prescribed measure of welfare. Specifically, when mandating disclosure results in a lower segmentation threshold λ^B, then both overall welfare and consumer surplus increase due to the mandate, such that there is no tension. However, when mandating disclosure results in a higher segmentation threshold λ^B, then overall welfare decreases but the effect on consumer surplus is inconclusive absent additional conditions. In the second period, consumer surplus will decrease, because the seller offers both segments of consumers a higher price. In the first period, however, while fewer consumers purchase the product – those who do may pay less for it, rendering the two-period effect on consumer surplus inconclusive.

The conclusion is that if mandating disclosure is welfare-enhancing, then it will also enhance the welfare of consumers. However, if mandating disclosure is welfare-decreasing, it may nonetheless increase the welfare of consumers. Note, however, that in any specific case, a careful examination may reveal that consumers will also be made worse by an inefficient mandate. Moreover, the consumers that may benefit from an inefficient mandate – are those who are willing to pay the most for the product, either those that are willing to purchase it despite the penalty of a higher future price or those the refrain from purchasing it strategically, and not because they truly value the product less. Often, in practice, these consumers are simply the wealthier, such that even under a policy that gives primacy to the welfare of consumers over efficiency, it may very well be unjustified to impose a mandate that is both inefficient and decreasing the welfare of weaker consumers, to the benefit of the wealthier ones.

5. Discussion

Unlike most of the typical techniques of price discrimination – that is based on inherent features of consumers, such as sex and age – behavior-based pricing does not necessarily lead to more efficient outcomes. On the one hand, as time progresses, customized pricing will lead to more consumers having access to the product or service sold. However, this comes at a cost, as both sellers and consumers are expected to change their behavior in the presence of behavior-based pricing. Some consumers will find it beneficial to strategically refrain from an otherwise efficient purchase, to secure lower future prices. Sellers, on the other hand, have contradicting incentives – increasing initial prices will increase the value of the information obtained from the decisions of consumers whether or not to purchase the product, while decreasing initial prices will make the product more appealing to strategic consumers. The result, however, is that fewer consumers will initially purchase the product. This cost should be taken into account by courts and policymakers when considering whether behavior-based pricing should be made (or interpreted as) illegal.

An alternative legal response is imposing a disclosure mandate on businesses that engage in behavior-based pricing. At the extreme, this could take the form of making algorithms transparent, possibly making the actual code of the algorithm public. However, this could also take the form of mandating businesses to publicly announce if their algorithms use prior behavior of consumers to set prices. Choosing the precise legal form of the mandate is outside the scope of this paper. Generally, however, it should balance between the legitimate business interests that firms have in keeping their algorithms secret with the effectiveness of any specific disclosure rule in informing consumers. At first glance, if *effective* disclosure is possible and not too burdensome on businesses, it may seem obviously desirable. Indeed, much of the prior literature treats algorithmic transparency as a desirable and relatively non-intrusive method of regulation. The common understanding is that when consumers misperceive elements of a pricing scheme and (1) it is not solved by market forces; (2) the method of disclosure effectively informs consumers; and (3) the direct costs to disclose are not exceedingly high, then mandating disclosure is viewed as desirable. This common wisdom does not hold when it comes to behavior-based pricing, because it fails to account for the possible inefficiencies caused by strategic consumers.

Given certain conditions as to the distribution of consumers in the market and the sensitivity of pricing decisions to it, as explored in section 4.2.4., mandating disclosure of behavior-based pricing could reduce the overall welfare (as well as the welfare of consumers). When consumers are made aware of the fact that a seller engages in behavior-based pricing, some of them react by refraining from a purchase that they would have made otherwise in order to secure lower future prices. Anticipating this resistance, the seller may decide to lower the price such that it will appeal to more hesitant consumers. This dynamic interaction determines whether disclosure will result in fewer or more consumers that purchase the product. If fewer consumers purchase the product, it follows that fewer consumers will purchase the product in the future as well, and the result is decreased efficiency because of the disclosure.

This finding presents a novel argument against mandatory disclosure, showing that ignorance could be bliss in the context of disclosure to consumers, even if the disclosure makes consumers rational, improves their decisions, and even if it does not impose any costs on businesses. The source of the inefficiency of disclosure lies in a type of collective action problem that has not been previously linked to disclosure mandates. Intuitively, a strategic consumer that does not purchase a product despite its low price, to secure a lower price in the future, is rational. Moreover, a single consumer that would change her behavior would not affect the pricing strategy of the seller. However, the *group* of consumers that strategically refrain from purchasing become indistinguishable from the group of consumers that did not purchase the product because it was too expensive to them. When facing these two indistinguishable groups, the seller is forced to set a single price to everyone in it. This drives the seller to set a higher price to everyone in this group relative to the price that would have been set if it was not for the strategic group's prior behavior. Therefore, strategic consumers make better choices than myopic consumers, potentially forcing the seller to lower the initial price, but by doing so they also generate a negative externality on their fellow consumers, by driving future prices up. Importantly, this collective action problem is created by the mandated disclosure, and would not have existed if sellers were permitted to remain vague as to their pricing schemes and consumers would remain blind to it.

While courts and policymakers may have in mind other valid justifications for promoting transparency, such as the value that consumers attach to their privacy or the disparate impact of price discrimination on various social groups, the results of this paper demonstrate the importance of tailoring the legal response to the particular properties of any business practice. In the context of algorithmic behavior-based pricing, it should be understood that disclosure and transparency may empower and improve the decisions of some consumers, but potentially through harming a significantly larger number of consumers, who are typically the less wealthy ones.

To be sure, these arguments rest on several assumptions that may not hold. The vast legal literature on disclosure mandates is focused on extensively scrutinizing each of these assumptions, and while this is outside the scope of this paper – algorithmic behavior-based pricing should be subjected to the same scrutiny. First, mandated disclosure may be irrelevant if behavior-based pricing is sufficiently easy to detect. While this may be the case in some settings, there are many others where firms are likely to find it relatively easy to conceal engaging in behavior-based pricing. For example, with airline tickets and ride-hailing services prices often change from minute to minute, and from one place to the other, making it difficult to detect behavior-based pricing without conducting a randomized controlled experiment. Second, disclosure may be ineffective in negating the myopia of consumers – they may either not pay attention to it or fail to account for it in their decisions. This is a general problem with disclosures, but as mentioned earlier – there is an important transition in the sense that more and more services are offered to consumers that act as intermediaries between consumers and the algorithms of sellers. These services, such as smart purchasing of airline tickets, empower consumers and level the algorithmic playing field without requiring any sophistication from the consumers themselves. Third, disclosure may itself be costly or burdensome. Most notably, forcing businesses to make their algorithms completely transparent may cause significant harm to the legitimate interests of businesses that developed the algorithms and protect them as trade secrets. However, this concern is more relevant to the decision of what type of disclosure to mandate. The algorithm could be sent discretely to a designated regulatory agency that will examine it without revealing any of its legitimate secrets. Moreover, algorithms could be made partially transparent, such that only the code that performs the regulated functions of it will be made public. Furthermore, disclosure may not involve making any part of the code public, but only to announce whether it performs certain functions that require disclosure.

More broadly than disclosure, an important issue is whether the findings in the paper are limited to markets controlled by a single seller, as the model assumes, or whether it extends to competitive settings. This is a general critique of models that are set in monopolistic markets, and arguably, the findings hold whenever a seller has some nonzero degree of market power, even if it is not a single seller. On the one hand, *perfect* competition would completely eliminate the type of behavior-based pricing considered in this paper, which is based on the willingness-to-pay of consumers, as it would eliminate most of the common practices of price discrimination – though not all of them.[30] However, markets are never perfectly competitive. For behavior-based pricing to be sustainable, it is not necessary that the seller is monopolistic, but only that it has some market power over its competitors. This may happen for any number of reasons, as long as purchasing from a certain seller entails some added value to consumers, such as better service, better shipping, loyalty programs (e.g., "Amazon Prime" membership) and so on. In terms of the model, the findings apply directly to any market where competition is imperfect, by

[30] Price discrimination practices can be crudely divided into two categories – those that discriminate based on the willingness-to-pay of consumers, and those that discriminate based on the varying costs of supplying a product or service to different consumers. Generally speaking, competition operates to curtail the former but encourage the latter. Practices such as offering customized insurance policies and customized loan rates are price discrimination practices that focus on the cost of supplying a product to different consumers. When a market is competitive, firms that do not engage in this type of price discrimination will experience adverse selection – only the riskier or more costly consumers will approach them, leading to losses and eventually market exit. On the other hand, price discrimination based on willingness-to-pay is hard to sustain in competition, because consumers that are offered the higher prices could go to another firm that will be willing to cut prices to snatch these consumers, and this competitive dynamic persists as long as prices are higher than the cost of production. Since the type of behavior-based pricing that this paper focuses on relates to the willingness-to-pay of consumers, perfect competition is expected to prevent it.

interpreting the value of the product to consumers as the *added* value from purchasing the product from a specific seller relative to the next best. In the business jargon, any level of product differentiation is sufficient for behavior-based pricing to be sustainable to some degree.

Another fundamental feature of the model is that it assumes that consumers and sellers interact only twice. While there are some markets in which this is a reasonable approximation for the frequency of repeated interaction, other markets exhibit significantly more frequent interactions between consumers and sellers. The frequent interactions could stem from the nature of the product, such as ride-hailing and flight tickets that many consumers purchase repeatedly over time. Moreover, it could stem from the nature of the seller, such as Amazon that sells an extremely wide variety of different products that can be priced based on a consumer's transactional history vis-à-vis the same seller, whether or not prior transactions involved the same product.[31] The model cannot apply directly to the case of multiple periods, but its results offer a grounded conjecture as to how behavior-based pricing is expected to play out when there are multiple periods. First, notice after one period the market is segmented into two parts, each offered a different price. Subsequently, in the following period, each of the segments could potentially be divided into two smaller segments. As the number of interactions increases, segments become smaller and smaller, wherein the extreme case where the number of periods goes to infinity, the seller gets closer to being able to engage in perfect price discrimination or perfectly customized pricing. However, as long as the number of periods is finite, this process is expected to progress slower as the number of future interactions is larger. Recall that in the two-period model, there were some strategic consumers who were better off refraining from an efficient purchase to secure lower prices. When the number of future periods increases, and as long as the discount factor is not exceedingly small, then more consumers will prefer to strategically avoid an initial purchase. The expected outcome is that in the initial periods, the seller would set some (possibly very low) price, but only those who value the product the most would purchase it. Therefore, the segmentation should not be perceived as repeatedly cutting the consumers roughly in halves, but rather as repeatedly cutting high-valuing consumers like cutting salami. While further study is required to explore this scenario, the expected dynamic emphasizes the lesson that was already revealed in the two-period model – repeated interactions will eventually generate many benefits, in the form of

[31] Behavior-based pricing does not require selling an identical product repeatedly. Purchasing some product could signal to the seller the willingness-to-pay of the consumer for a different product, as long as the willingness-to-pay for the two products are correlated, although this is outside the scope of the model in this paper.

almost-perfect price discrimination, but at a great deal of cost, since learning will be gradual. In potentially many of the first periods, very few consumers will purchase the product.

Behavior-based pricing, both in the two-period model and more so when interactions repeat more than twice, give rise to an interesting notion that I term "consumer reputation". [32] Just as businesses build their reputation through reviews by consumers, so can consumers be seen as building their reputation vis-à-vis businesses. One prominent manifestation of consumer reputation is credit scores, which tell different lenders what is the cost of extending a loan to a specific borrower. Just as the reputation of a business acts to signal to consumers what is the value of the product or service offered by businesses (e.g., a hotel), so could the reputation of a consumer signal to businesses what is their willingness-to-pay for a product. The transactional history of consumers – which products they purchased and at what price – could serve as the information underlying the reputation of consumers. Just as consumers can use the reputation of hotels to decide which hotel to choose, so can sellers use the reputation of consumers to decide what prices they should offer them. The analytical analogy between reputation and behavior-based pricing should inform the literature on reputation of the potential perils of it. While the reputation of businesses is often celebrated as an efficient tool for self-regulation by the market, this is not the case when the consumers are the subject of reputation. As shown, the prospect of establishing a reputation could distort the incentives of consumers to consume, even if it eventually results in more efficient pricing. In principle, there is no reason to think that these inefficiencies cannot extend to businesses, which, by analogy, may over-invest in building their reputation. In practice, this does not seem like a serious concern, but it is important to understand why. Businesses are assumed to potentially live 'forever' and interact with numerous consumers – giving rise to the intuition articulated earlier, that when interactions are numerous, the long-term efficiency of reputation may exceed the costs of accumulating it. Whether or not this is a valid conjecture, it certainly does not hold with regard to the reputation of consumers, who often interact just several times with any single seller. While a rigorous analysis is outside the scope of the paper, studies on reputation should be less quick to assume that there are no concerns from the process of building a reputation, as this paper provides an example to a case where the social costs of creating reputation may exceed its benefits. For example, imagine that Google was to use its web browser Chrome to collect all data on purchases that consumers make online, as an intermediary, and sell the information to interested sellers, who may be interacting with a consumer for the first time while already in possession

[32] I thank Louis Kaplow for suggesting this point.

of her entire history as a consumer. Should we be concerned about such a reputational mechanism, or embrace it similar to embracing the reputation of businesses?

Before concluding, it is important to address the issue of the potential disparate impact of behavior-based pricing on various social groups. Throughout the paper, we assumed that consumers are homogenous in all but their willingness-to-pay for a product. However, willingness-to-pay may consist of various factors,[33] among them the wealth of a person. People may have lower willingness-to-pay simply because they are less wealthy, and people who are less wealthy are, unfortunately, disproportionally black, Hispanic, and women. Arguably, using the race of a person as a direct input in a pricing algorithm is different from using their willingness-to-pay, which is only statistically correlated with their race. However, unlike in the context of labor markets, in consumer markets, neither practices are affirmatively illegal (as long as they do not lessen competition). The question of whether behavior-based pricing should be prohibited because of its disparate impact is an important one, but outside the reach of this paper. Note, however, that in contrast to the lending markets, for example – where it is well-documented that black people are offered higher interest rates – this is unlikely to be the case with behavior-based pricing. Generally speaking, behavior-based pricing – as well as any type of price discrimination that proxies the willingness-to-pay rather than the cost of producing a product – lead to *lower* prices being offered to consumers with low willingness-to-pay. If these people are disproportionally black, or Hispanic, then they may be those that benefit the most from behavior-based pricing, while the wealthy are the ones that will be offered higher prices. For this reason, it seems that concerns about equality should not lead to prohibiting behavior-based pricing if it is welfare-enhancing.

6. Conclusion

Recent years have seen a surge in the legal scholarship regarding the desirable legal response to the use of algorithms by businesses to price discriminate. However, this scholarship tends to bundle together different types of discriminatory pricing practices, ignoring paramount differences between them. Behavior-based pricing is an important business practice that has been resurrected in the algorithmic age and currently stands in the technological frontier of algorithmic pricing. In addition to the importance of studying it as a central phenomenon that may call for tailored regulation, it demonstrates why the legal scholarship should be more attentive to the economic properties of different business practices when

[33] See footnote 5, *supra*, for a detailed account on the construction of the willingness-to-pay of a consumer for a product or service.

prescribing the desirable legal response. When it comes to behavior-based pricing, price discrimination may be welfare-reducing, unlike classic techniques of price discrimination. Moreover, even if behavior-based pricing is welfare-enhancing – mandating the disclosure of it may result in adverse outcomes, in contrast to the hopes that are often expressed by advocates of algorithmic transparency.

References

- Acquisti, A., & Varian, H. R. 2005. Conditioning prices on purchase history. *Marketing Science*, 24(3):367-381.
- Bar-Gill, O. 2020. Consumer Misperception in a Hotelling Model: With and Without Price Discrimination. *Forthcoming, Journal of Institutional and Theoretical Economics*.
- Bar-Gill, O. 2019. Algorithmic Price Discrimination: When Demand Is a Function of Both Preferences and (Mis)Perceptions. *University of Chicago Law Review*, 86:217-254.
- Benkler, Y. 2019. Don't let industry write the rules for AI. *Nature*, 569(7755): 161-161.
- Borenstein, S. (1992). The evolution of US airline competition. *Journal of Economic perspectives*, 6(2):45-73.
- Chen, L., Mislove, A., & Wilson, C. 2015. Peeking beneath the hood of Uber. In *Proceedings of the 2015 internet measurement conference*, 495-508.
- Chen, Y., & Zhang, Z. J. 2009. Dynamic targeted pricing with strategic consumers. *International Journal of Industrial Organization*, 27(1):43-50.
- Cui, R., Li, J., Li, M., & Yu, L. 2020. Wholesale Price Discrimination in Global Sourcing. *Forthcoming, Manufacturing & Service Operations Management*.
- Desai, D. R., & Kroll, J. A. 2017. Trust but verify: A guide to algorithms and the law. *Harv. JL & Tech.*, 31:1-64.
- Fudenberg, D., & Villas-Boas, J. M. 2006. Behavior-based price discrimination and customer recognition. *Handbook on economics and information systems*, 1:377-436.
- Hart, O. D., & Tirole, J. 1988. Contract renegotiation and Coasian dynamics. *The Review of Economic Studies*, 55(4):509-540.
- Kleinberg, J., Ludwig, J., Mullainathan, S., & Rambachan, A. 2018. Algorithmic fairness. *AEA papers and proceedings*, 108:22-27.
- Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. 2018. Discrimination in the Age of Algorithms. *Journal of Legal Analysis*, 10:113-174.

- Kroll, J. A., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., & Yu, H. 2016. Accountable algorithms. *U. Pa. L. Rev.*, 165:633-705.
- Li, J., Granados, N., & Netessine, S. 2014. Are consumers strategic? Structural estimation from the air-travel industry. *Management Science*, 60(9):2114-2137.
- Liu, Q., & Van Ryzin, G. J. 2008. Strategic capacity rationing to induce early purchases. *Management Science*, 54(6):1115-1131.
- Phillips, R. L., & Solutions, N. 2012. Customized pricing. *The Oxford Handbook of Pricing Management, Oxford University Press, Oxford*, 465-490.
- Rothschild, M. 1974. A two-armed bandit theory of market pricing. *Journal of Economic Theory*, 9(2):185-202.
- Su, X. 2007. Intertemporal pricing with strategic customer behavior. *Management Science*, 53(5):726-741.
- Taylor, C. R. 2004. Consumer privacy and the market for customer information. *RAND Journal of Economics*, 631-650.
- Yang, C., & Dobbie, W. 2019. Equal Protection Under Algorithms: A New Statistical and Legal Framework. *Available at SSRN 3462379*.