Dispersion, Discrimination, and the Price of Your Pickup^{*}

Stephan Sagl[†]

May 7, 2024

Abstract

Using repeat purchase data on pickup trucks, I establish that the same consumers pay persistently high or persistently low prices across vehicle purchases. Less than 1% of this persistence can be explained by demographics. This result suggests that dealers use consumer information beyond coarse demographics to personalize prices. Using a novel discrete choice model with personalized pricing, I study the role of consumer information firms use for pricing in the welfare effects of price discrimination. To do so, I overcome a common problem in settings with transaction data: personalized prices of non-chosen alternatives are unobservable. I solve this problem by recovering unobserved personalized prices and consumer-specific price sensitivity from the observed transaction price via firms first-order conditions. I simulate two counterfactuals: uniform pricing and price discrimination based on coarse demographic groups. Compared to uniform pricing, personalized pricing increases profits and total welfare but, on average, harms consumers. On the other hand, compared to uniform pricing, price discrimination based only on demographics is not profitable. This highlights the importance of the amount of consumer information firms can use for pricing in the welfare effects of price discrimination.

JEL Codes: D43, L13, L62

^{*}I thank Paul Grieco, Charles Murry, Joris Pinkse, Mark Roberts, Connor Ryan, Bradley Setzler, and Karl Schurter for invaluable guidance. I thank all Penn State IO brown bag participants for helpful comments. This paper relies on data on sticker prices and manufacturer-installed options. I thank BlackBook and a leading online used car retailer for providing me with these data.

[†]Department of Economics, The Pennsylvania State University, stephan.sagl@psu.edu.

1 Introduction

Price discrimination is a pervasive, yet controversial feature of the US automobile market. However, whether price discrimination benefits firms or consumers is theoretically ambiguous. In particular, price discrimination can benefit firms, some consumers, or all consumers. The direction of the welfare effects depends, among others, on two factors: the precise form of consumer heterogeneity and the consumer information firms use for pricing. Academic literature has focused on studying price discrimination by assuming firms leverage coarse consumer information such as demographic groups for pricing and found generally small or insignificant effects (e.g. Langer, 2016; D'Haultfoeuille, Durrmeyer and Février, 2019). In this paper, I provide novel evidence that firms use much more granular, individual-level consumer information in their pricing decisions to personalize prices and develop a novel framework to study the role of consumer information in the welfare effects of price discrimination.¹

Using data on pickup truck purchases in Texas, I make three main contributions. First, leveraging data on consumers who bought trucks repeatedly and fully controlling for product heterogeneity, I establish that firms set personalized prices. After controlling for pickup heterogeneity, around 25% of price dispersion is due to price discrimination based on consumer characteristics. However, only 0.5% of price dispersion is due to demographics observable to the econometrician. Second, I solve a key problem in demand estimation models in which prices are personalized or consumer-specific: the personalized prices of non-chosen alternatives can never be observed. I overcome this problem by developing and estimating a novel discrete choice model of supply and demand that recovers unobserved personalized prices and consumer-specific price sensitivity from the observed transaction price exploiting firms' pricing first-order conditions. I do not make parametric assumptions on the joint

¹Literature commonly refers to third-degree price discrimination based on a rich set of consumer characteristics as *personalized prices* (e.g. Buchholz et al., 2022; Dubé and Misra, 2023).

distribution of price sensitivities and demographics. Price sensitivities can depend arbitrarily on consumer demographics. Third, I use this model to study the role of information in the welfare effects of price discrimination. I find that consumer-specific, personalized prices raise profits by around 19% relative to uniform pricing, increase total surplus, and harm consumers on average. However, how much consumer information firms use for pricing is crucial for the direction of the welfare effects. Restricting firms to only use information on coarse demographics in their pricing decisions lowers profits by 0.3% relative to uniform pricing.

I establish these findings by leveraging new data on pickup truck purchases linked to both the actual buyers and the actual pickup truck characteristics including installed optional equipment. Using data on sticker prices inclusive of options and delivery fees, I find that different consumers pay different prices for the same good. After controlling for optional equipment and delivery fees, the average within-model standard deviation is around \$4,100. Leveraging data on repeat purchases, I use Abowd, Kramarz and Margolis (1999)'s two-way fixed effects estimator to establish that the same consumers pay persistently high or persistently low prices across vehicle purchases, controlling for product characteristics. Around 25% of price variation after controlling for product heterogeneity is due to price discrimination. This shows that persistent consumer preferences are known to dealers and used when making price offers.

Further decomposing the variance of estimated consumer effects, I find that even a wide array of consumer demographics, including income, only explains 0.5% of price variation after controlling for product heterogeneity. This suggests that dealers have a large amount of consumer information beyond coarse demographics which they may glean from test drives, conversations, credit scores, etc. This finding is reminiscent of the use of personalized pricing - pricing based on a rich set of consumer features in the marketing literature (e.g. Dubé and Misra, 2023).

To understand the importance of consumer information available to firms in

the welfare effects of price discrimination, I estimate a structural model of supply and demand where dealers observe key consumer preferences and price discriminate accordingly. The model flexibly incorporates consumer heterogeneity via observed demographics and unobserved tastes, including price sensitivity. The key challenge complicating estimation is that the personalized prices of non-chosen alternatives can never be observed. To overcome this problem, I use firms' pricing first-order conditions to recover unobserved prices of non-chosen alternatives and price sensitivities directly from the observed transaction prices. This allows me to learn the joint distribution of price sensitivities and demographics for all consumers in the market while avoiding to make parametric assumptions on the joint distribution.

I use this model to evaluate the role of consumer information in the welfare effects of price discrimination. First, I consider the elimination of price discrimination. It is well known that when firms compete, price discrimination may increase or decrease consumer surplus, depending on the precise form of consumer heterogeneity. In addition, price discrimination may also decrease firms' profits (Corts, 1998). I find that all producers benefit from price discrimination in the pickup truck market, and consumers are harmed on average.

Secondly, I illustrate the importance of consumer information firms use for pricing for the direction of the welfare effects of price discrimination. To do this, I consider a counterfactual where firms price only based on observable consumer demographics but do not use additional information about consumers. In this case, all firms lose, and industry profits are lower by around 0.3% relative to uniform pricing. This demonstrates that the welfare effects of price discrimination are highly sensitive to the amount of consumer information available to firms for pricing. Furthermore, these results show that dealers use much richer information than gender and ethnicity to price discriminate. In particular, most of the price dispersion and profits stem from dealers leveraging the information they acquire through direct interactions with consumers rather than coarse demographics.

Why does price discrimination not lead to lower prices in the pickup truck market? Corts (1998) showed that if firms differ sufficiently in their ranking of consumers in terms of demand elasticities, best response asymmetries can induce discriminatory prices below uniform prices. My estimates imply that the drivers of price dispersion in the pickup truck market are differences in price sensitivities across consumers rather than demographics. Because all dealers rank price-insensitive consumers higher and this channel dominates, there is no sufficient asymmetry in best responses. Thus, price discrimination did not lead to lower prices in the pickup truck market.

Previous research has studied the average effects of demographics on price (e.g. Goldberg, 1996; Harless and Hoffer, 2002; Chandra, Gulati and Sallee, 2017). My estimates leveraging repeat purchases of individual consumers complement and expand on these findings. While I find mean differences in price paid across demographics, my estimates demonstrate that price dispersion is almost entirely driven by unobservable heterogeneity or soft information. These results also complement recent work on price discrimination in the adjacent auto loan intermediation market finding that dealers mainly levarage soft information to price discriminate (Grunewald et al., 2023).

This paper also offers a contrasting perspective on recent structural work on price discrimination in the automobile market (Langer, 2016; D'Haultfoeuille, Durrmeyer and Février, 2019). Existing literature studying the retail market has found small effects of price discrimination on firms' profits. My results offer a stark contrast to these findings. Dealers engage in price discrimination, and it is highly profitable to them. However, in line with earlier structural research, I find that price discrimination based on protected classes or observed demographics only leads to small changes in profits relative to uniform pricing.

My work also contributes to the literature on the estimation of discrete choice demand models when either prices or quantities are unobserved (e.g. Thomadsen, 2005; Coşar, Grieco and Tintelnot, 2015; Grigolon, Jacobi and Sovinsky, 2020). Methodologically closest to my approach is D'Haultfoeuille, Durrmeyer and Février (2019), who develop a model of supply and demand with unobserved price discrimination based on coarse demographic groups. I contribute to this literature by building and estimating an equilibrium model of supply and demand where firms set *individual* (*personalized*) rather than group-specific prices when researchers can only observe the realized transaction price.

I also contribute to the empirical literature on the welfare effects of personalized pricing. A growing subset of this literature has recently focused on personalized pricing in online markets (Dubé and Misra, 2023; Shiller, 2020, 2022). However, there has been less work on offline markets (Waldfogel, 2015; Crawford, Pavanini and Schivardi, 2018; Buchholz et al., 2022). I contribute to this literature by studying an offline, multi-firm, oligopolistic market under a baseline scenario of personalized prices using an equilibrium model of demand and supply. I show that personalized prices benefit firms, harm consumers on average, and increase total welfare. However, with different information structures welfare effects flip.

I organize the remainder of this paper as follows: Section 2 describes the data used in this paper and its sources. Section 3 presents evidence for personalized pricing in the market for pickup trucks. Section 4 introduces the equilibrium model of supply and demand. Section 5 presents the identification argument. Section 6 presents the empirical strategy. Section 7 presents the results. Section 8 studies the role of consumer information in the welfare effects of price discrimination, and section 9 concludes.

2 Pickup Trucks & Data in Texas

This section describes the market for pickup trucks in Texas and the data I use for this paper.

2.1 Industry Background

The automobile market is one of the US's most important consumer goods markets. A defining feature distinguishing the US automobile market from other countries' markets is the popularity of personal-use pickup trucks. Of the around 17 million new light-duty vehicles sold across the US in 2018 (National Automobile Dealers Association, 2019*b*), nearly 2.4 million, or about 14%, were full-size pickup trucks (Drury and Caldwell, 2020). Remarkably, pickup trucks made up three of the top three most popular vehicles in the US in 2020, with the Ford F-Series claiming the top spot as America's most popular vehicle for nearly four decades (Wayland, 2021).

While pickup trucks are popular nationwide, pickup trucks are particularly popular in Texas. Based on my data, between 2011 and 2019, around 22% of all new vehicle sales in Texas were pickup trucks. Focusing on Texas is important because the pickup truck market in Texas is also nationally significant. According to data from Experian, Texas was the biggest market for pickup trucks in the US in the first quarter of 2019 by volume, accounting for around 13% of all pickup truck registrations in the US (Miller, 2019). The vast majority of new pickup trucks in Texas, with limited exceptions, sell through franchised dealerships.²

In 2018, there were around 1,200 new vehicle dealerships in Texas (National Automobile Dealers Association, 2019a). New automobile dealerships in Texas are commonly franchised with a single manufacturer, but sometimes larger dealerships

²Exceptions are electric vehicle manufacturers like Rivian or Tesla, who sell pickup trucks directly to consumers. This is a relatively recent phenomenon.

are franchised with more than one manufacturer.³ Dealerships only have access to inventory of new pickup trucks from manufacturers they are franchised with. While franchise regulations are complex and out of the scope of this overview, Murry and Schneider (2016) argue that current regulations prohibit manufacturers from terminating dealerships in all but exceptional cases, prevent manufacturers from wholesale price discrimination across dealerships, and allow dealerships to set automobile prices freely.

The sale of a pickup truck commonly proceeds in a standardized way. Traditionally, potential customers would enter a dealership's lot looking for the pickup truck they are interested in.⁴ In most cases, a salesperson then approaches the customer, informing the customer about specific vehicle features, asking questions about the customer's background, and eventually offering a test drive.⁵ If the customer expresses interest in a test drive, she is usually allowed to test drive the pickup truck for a limited time. After the test drive, the salesperson asks the customer to discuss the potential purchase terms. Discussing the terms is commonly the most time-consuming part of buying a new pickup truck, with US dealership visits averaging close to 3 hours in 2017 (Cox Automotive, 2019). During the negotiations phase, the salesperson will make multiple offers and will sometimes leave to "consult" with their manager behind closed doors.⁶ At the end of the negotiations phase, the salesperson will typically state that the current price is the best price they can offer. If the consumer decides to buy, the salesperson usually earns a significant commission on the sale of the pickup

 $^{^3{\}rm For}$ example, Vandergriff owns an Acura, a Chevrolet, a Honda, a Hyundai, and a Toyota dealership in Arlington, TX, located just a few feet apart next to Interstate 20.

⁴Although most new vehicle dealerships in the US offer appointments with salespersons on request, walk-ins remain popular: according to Cox Automotive (2019), around 50% of customers are walk-in. If a consumer schedules such an appointment, she usually enters the dealership to ask for the salesperson instead and proceeds from there.

⁵These interactions allow the salesperson to learn more about their consumers. Common questions are the number of children in the household, the vehicle's intended use (commute/pleasure), etc.

 $^{^{6}\}mathrm{I}$ use "consult" because often there is no consultation, and the sales person only makes the customer wait.

truck.⁷

The distinguishing feature of the pickup truck and automobile market relative to other markets is the lengthy interactions between consumers and dealers. Literature has found that quoted prices during negotiations correlate with demographics, even when auditors follow the same negotiation protocol (Ayres and Siegelman, 1995). Zettelmeyer, Scott Morton and Silva-Risso (2006) further suggest that the negotiations phase may serve as a way for dealers to infer consumers' willingness to pay and to price discriminate. This suggests that sales personnel combine the information gleaned from test drives, credit scores, and conversations with observable consumer demographics to learn about customers' willingness to pay.

2.2 Data Sources

I rely on a combination of different data sets. My main sources are *transaction-level* registrations data from the Texas Department of Motor Vehicles (TX DMV), detailed data on buyer characteristics provided by *Infutor*, as well as data on pickup truck characteristics up to the manufacturer installed options obtained from leading industry sources.⁸

TX DMV Registrations data. My primary data set consists of individual transaction-level data on all vehicle purchases for the state of Texas obtained from the Texas Department of Motor Vehicles from January 1, 2011, to October 31, 2019. Each of the observations corresponds to a unique vehicle sale. I observe the make/model of the vehicle, the month of purchase, the vehicle identification number (VIN), the transaction price, the name of the dealership selling the vehicle, its five-digit zip code, and the five-digit zip code of the buyer. I limit the sample to pickup trucks and

⁷Literature reports that these commissions are commonly around 20 to 30 percent of the profit margin of the dealership (Murry and Schneider, 2016, p. 345).

⁸I combine data from multiple industry sources to construct my pickup truck characteristics data. Only BlackBook has agreed to have its identity disclosed. Other data sources include a leading online used car retailer.

exclude models mostly used commercially.⁹ Additionally, I geo-code buyer and dealer locations to obtain purchase distances. These data consist of 1,433,657 transactions spanning the whole state and nine years.

infutor Consumer data. For consumer demographics, I rely on data collected by *infutor Data Solutions*, henceforth infutor. The infutor data recently became popular for research requiring tracking individuals' US addresses (e.g. Bernstein et al., 2022; Diamond, McQuade and Qian, 2019) because it provides persistent identifiers for and detailed demographic data on individuals. It is highly representative of the overall US adult population and covers 78% of the overall adult US population (Bernstein et al., 2022, p. 7). I impute the ethnicity of individuals using NamePrism (Ye et al., 2017; Ye and Skiena, 2019).¹⁰ While previous studies exploited infutor's data on housing, I leverage infutor's auto data. According to Infutor Data Solutions (2023), infutor's auto data contains nearly 200 million vehicle owner records built from information on sales records, service records, and auto repairs, among others. Besides information on the owner, the auto data contains basic information about the vehicle plus, importantly, the complete 17-digit vehicle identification number (VIN).¹¹ I use the vehicle identification number to merge the demographic data from infutor with my transaction-level data. Combining these data allows me to (i) learn the identity and, thus, the demographics of purchasers; and (ii) track automobile purchases of individuals over time.

Sticker Price Options data. In the US, by law, manufacturers must affix a label stating, among other information about the vehicle, the manufacturer's suggested retail price of the automobile and all installed optional equipment to every new

 $^{^{9}\}mathrm{I}$ exclude trucks such as, e.g., the Ford F-350 and F-450 because these are rarely used as personal vehicles.

¹⁰NamePrism is a tool to classify ethnicity and nationality based on names widely employed in academic research.

¹¹Note that infutor's auto data comes standard with a masked 10-digit VIN. The full 17-digit VIN has to be requested explicitly.

vehicle before delivery (15 U.S.C. §1232). These labels are called *window stickers*.¹² Information on the window sticker allows consumers to learn the exact specifications of the pickup truck they are about to purchase. I obtain data on these window stickers from two industry sources. My primary data source is *Black Book*. Black Book is an industry-leading data provider for automobile dealerships in the US. Secondly, I supplement these data with data from a leading online used car retailer. In total, I obtain data on approximately 1.65 million window stickers.¹³ Again, I merge these data with my transaction data based on the vehicle identification number.

Merging all three data sets, I end up with 76,675 matched observations. Table 1 documents differences in the purchase price, base price, price of the optional equipment, and income across demographic groups. On average, female and Hispanic consumers choose less expensive truck models and trims and spend less on optional equipment. These demographic groups also have lower incomes.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Female	Male	Non-Hispanic	Hispanic	income	income	income
					\leq \$50k	\in (\$50k,\$100k]	>\$100k
Purchase Price	$36,\!140.54$	$37,\!650.31$	$37,\!534.96$	36,907.46	$36,\!683.47$	37,051.95	38,465.21
	(7,951.64)	(8, 439.60)	(8,344.76)	(8,375.34)	(8, 282.32)	(8,227.94)	(8,637.09)
Base Price	37.322.01	39.286.77	39.280.95	38.031.22	37.716.22	38.648.64	40.204.43
	(7,943.15)	(8, 123.13)	(8, 119.83)	(8,070.98)	(8,069.30)	(8,033.74)	(8,220.08)
Options Price	$5\ 116\ 72$	5 660 08	5 582 44	$5\ 465\ 05$	5 302 06	5 521 23	5 765 59
options i nee	(3,607.76)	(3,730.93)	(3,748.43)	(3,635.27)	(3,577.39)	(3,651.53)	(3,934.78)
Ŧ		~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~		00 F 01 01			
Income	74,467.07	82,542.82	87,809.69	66,791.61			
	(28, 848.58)	(32,013.39)	(31, 527.39)	(26, 473.93)			
Ν	$16,\!475$	60,200	51,131	$25,\!544$	12,463	46,099	18,113

Table 1: Summary Statistics across Demographic Groups

My data offers several advantages. First, it contains data on the actual pickup truck buyers, their demographics, and complete information on pickup truck characteristics.

 $^{^{12}\}mathrm{Window}$ stickers are sometimes called Monroney labels.

¹³These are window stickers for all types of vehicles in the transaction level data, not just pickup trucks. Each window sticker matches one vehicle in the transaction data.

Second, because infutor tracks consumers over time, I can construct a panel of repeat purchases for a subset of consumers. Across all years, I identify 2,640 repeat pickup truck purchases by the same consumers. Because the repeat purchase data set is relatively small, however, I do not use it for structural estimation. In particular, I construct two different data sets from the 76,675 matched observations. First, the *repeat purchase data set* which covers 2011 to 2019. Second, the *microdata for demand estimation* I use to estimate my structural model covering 2016 – 2019. The second data set does not cover all years because the merge rate before 2016 is significantly worse than for 2016 – 2019. The drop in the merge rate is due to one of my data sources only becoming an increasingly important player in the used car retailing business after 2015, significantly reducing the amount of sticker data I can match with the transaction level data.

Lastly, a commonly encountered problem in structural work on the US automobile market is that estimating a model at the trim level is impossible because the product space would be prohibitively large (e.g. Berry, Levinsohn and Pakes, 1995; Langer, 2016; Murry, 2017). For this reason, I normalize pickup trucks to the base trim in the microdata for demand estimation data set with no installed options. To do so, I first subtract the price of the installed options from the transaction price. I then remove the price effect of the trim by running a regression of transaction price net of options on fixed effects for trims for each pickup truck model separately. The normalized price is then the average price of the base trim plus the individual specific residual from this regression. This preserves the price dispersion and is similar to the homogenization of bids commonly used in the auction literature. Note that I do not need this normalization in the repeat purchase data set. The documentation of the price dispersion in section 3 is unaffected.

3 Documenting Price Discrimination

In this section, I present evidence that motivates a model in which firms set individual consumer-specific prices, called *personalized prices*, based on rich consumer information. In particular, this section documents the following facts about the pickup truck market. First, different consumers pay vastly different prices for the same pickup truck model. Second, the same consumers pay persistently low or persistently high prices across vehicle purchases. Third, demographics only explain a tiny fraction of price dispersion suggesting that firms mostly use soft information in pricing.

Fact 1: Different consumers pay different prices for the same pickup truck

I start by documenting the first fact: different consumers pay vastly different prices for the same pickup truck model. Prices of a pickup truck model can vary because of differences in product offerings or price discrimination. I first establish that differences in product offerings are important for observed price dispersion but that they cannot fully rationalize it. While all pickup truck models in my data exhibit roughly the same patterns, focusing on a single model is easier because it eliminates between-model heterogeneity. In the following, I present statistics for all 2018 Chevrolet Silverado 1500s sold in Texas in 2018.¹⁴

Pickup trucks are heterogeneous products, even conditional on model and trim. The problem of product heterogeneity is well known but widely ignored in the literature because of data limitations, e.g., Berry, Levinsohn and Pakes (1995). Consumers can customize pickup truck models by selecting from various options. For the 2018 Chevrolet Silverado 1500, the average consumer chooses 16 out of the 200 available options in 2018.¹⁵

 $^{^{14}}$ I use the Chevrolet Silverado 1500 as an example because it was the second most popular truck in the US in 2018 (Wardlaw, 2019).

¹⁵Of the 200 options, some seem to be the same but have different option codes. I therefore report them as different options.



Figure 1: Distributions of optional equipment choices. Left panel plots frequency of specific options installed on share of all 2018 Chevrolet Silverado 1500s. Right panel plots distribution of the price of options.

The left panel in figure 1 shows that most options are only installed on a handful of pickup trucks. Relatively few options are installed on the majority of pickup trucks. These popular options, like 10-way adjustable driver seats, remote start ability, or electric rear windows, usually come in packages. Rarer options like spot lamps, upgraded spare wheels, or suspension options do not occur more than a few times in the data. However, less common options are, on average, more expensive.

The right panel in figure 1 documents the distribution of options prices. The mean total price for all installed options on a 2018 Chevrolet Silverado 1500 is around \$6,500. While most options cost below \$1,000, with an average of around \$825, many installed options can get very expensive, with the maximum exceeding \$7,000.

While options are important factors that drive the dispersion of transaction prices, figure 2 provides evidence that options alone are insufficient to explain the price variation. Figure 2 plots the distribution of two different margins for the 2018



Figure 2: Density of Margins when controlling for Options (red) and when not controlling for Options (blue)

Chevrolet Silverado 1500: the transaction price minus the sticker price (red),¹⁶ and the transaction price minus the base MSRP at the trim level (blue) which does not include any manufacturer-installed options. If manufacturer-installed options indeed entirely drive the price dispersion, then the transaction price minus the sticker price (red) distribution would degenerate to a single point. Looking at figure 2, this is not the case. While often proposed as one reason for price dispersion, accounting for manufacturer-installed options even slightly increases the standard deviation by around \$300 to \$4,200. Differences in transaction prices do not only stem from people opting for different levels of packages or manufacturer-installed options. Additionally, because the sticker price fully controls for differences in product offerings, the remaining price differences are solely due to price discrimination.¹⁷

 $^{^{16}{\}rm The}$ sticker price is the manufacturer's suggested retail price including all installed options. It is not binding for dealers.

¹⁷Note that price discrimination alone does not imply that prices are personalized.

		Price		Price - Sticker
	(1)	(2)	(3)	(4)
Female	$295.81 \\ (244.79)$		262.61 (246.34)	$729.67^{***} \\ (198.09)$
Hispanic		417.49^{*} (235.23)	399.18^{*} (236.91)	$\begin{array}{c} 603.40^{***} \\ (193.54) \end{array}$
Income (10,000s)	46.18 (33.33)	59.62^{*} (34.05)	61.60^{*} (34.09)	-20.01 (28.35)
Controls:	,	,	,	,
Year-Model-Trim FE	\checkmark	\checkmark	\checkmark	\checkmark
NISA FE	√ 2.280	√ 2.280	√ 2.280	√ 2.280
R^2	0.86	0.86	0.86	0.72

Table 2: Sticker prices uncover mean differences in prices paid across demographics

Note: Number of observations smaller than 2,640 because some observations miss income. Standard errors in parentheses

* p < 0.10, ** p < 0.05, *** p < 0.01

Fact 2: Prices are individual consumer-specific

Next, I document that pickup truck prices are individual consumer-specific, and that consumers pay persistently high or persistently low prices across vehicle purchases. To do so, I decompose the variance of the difference between the transaction prices and sticker prices for all pickup truck models and all years into a consumer and product component. This allows me to explore how much of the variation in transaction prices is due to price discrimination on consumer characteristics.

The repeat purchase data enable me to follow a subset of consumers across purchases. I first show that the repeat purchase data are consistent with existing literature studying price discrimination in the US automobile market. Literature has repeatedly found mostly statistically insignificant or small effects when regressing transaction prices on a broad set of consumer demographics and vehicle fixed effects even on the model-trim level. Table 2 shows that my repeat purchase data are consistent with these findings. Even when controlling for trim, all demographic variables fail to achieve significance at the 5% level. The lack of statistical significance arises because demographics correlate with trim levels and optional equipment. Column (4) shows that appropriately controlling for trim and optional equipment uncovers correlations between consumer demographics and prices paid.

Because I observe consumers and pickup trucks repeatedly, I can decompose the variance following Abowd, Kramarz and Margolis (1999), henceforth AKM. This is a novel approach in the automobile market because researchers usually only have access to cross-sectional data. Following AKM, I estimate the following equation using repeat purchases of the same customers¹⁸

(Transaction Price - Sticker)_{ijt} =
$$z'_{it}\beta + \eta_i + \phi_{j(i,t)} + \epsilon_{it}$$
 (1)

where (Transaction Price - Sticker)_{*ijt*} is the difference between transaction price and sticker price for consumer *i* purchasing pickup truck *j* at time *t*, z_{it} are timevarying characteristics, η_i is a consumer-specific fixed effect, and $\phi_{j(i,t)}$ are pickup truck year-model-trim fixed effects.

AKM's variance decomposition allows me to distinguish between variation in transaction prices due to differences in product offerings and price discrimination. I estimate varying specifications of equation (1) and present the results in table 3.

My main result is that time-varying covariates do not contribute to the price dispersion but that time-invariant, consumer-specific fixed effects η_i account for around 25% of the variation in transaction price minus sticker price (column 4).¹⁹ This

¹⁸Note that in the AKM framework, the dependent variable is usually put in logs. I do not follow this because logit pricing first-order conditions in Nash-Bertrand models are not log-separable.

¹⁹Note that the sum of the fractions of variance explained do not have to sum to one because of the unreported covariances. This is a well known fact in the labor literature.

Table 3:	Decomposition	of the	variance	following	Abowd,	Kramarz	and	Margolis
(1999)								

dependent variable: (price - sticker)_{ijt}	(1)	(2)	(3)	(4)
Fraction of variance explained:				
$z_{it}^{\prime}eta$.00	.00	.00
individual FE η_i	.25	.25	.25	.25
pickup truck FE $\phi_{j(i,t)}$.72	.72	.72	.72
match heterogeneity ϵ_{it}	.10	.10	.10	.10
Time varying characteristics:				
Purchase Distance (mi)		-15.35 (16.99)		-15.19 (16.98)
Financed			$789.90 \\ (567.61)$	$786.41 \\ (567.67)$
Ν	2,640	2,640	2,640	2,640
Standard errors in parentheses				

* p < 0.10, ** p < 0.05, *** p < 0.01

suggests that dealers have access to large amounts of information about persistent consumer characteristics they use for pricing. It also implies that consumers pay persistently high or persistently low prices across vehicle purchases. In other words, prices follow people across pickup truck purchases. This is a novel result not shown before in literature on the US automobile market.

The remaining columns investigate whether other factors also explain the variation in the transaction price. Column 2 explores whether search costs could be the driving factor of price variance by controlling for purchase distance. While search costs might be otherwise significant, variation in purchase distance does not explain the variation in the transaction price. The coefficient on purchase distance has the expected sign, but it is not significant even at the 10% level. The same is true for the financing decision of the consumer, and jointly controlling for both financing and purchase distance. However, even including all these controls jointly does not explain the variation in transaction price since the contribution to price variation is estimated to be approximately 0%.

Fact 3: Coarse demographics only explain a tiny fraction of price dispersion

Lastly, I establish that coarse demographics only explain a tiny fraction of price dispersion. To understand the importance of price discrimination based on demographics versus information firms can learn through interactions with consumers, I project the estimated consumer-specific effects on demographic controls: gender, ethnicity, the mean income across purchases for customer i, and the mean purchase distance across purchases for customer i. I also add fixed effects for Metropolitan Statistical Areas (MSAs) that control for differences in competition and population density.

The results in table 4 show that the consumer-specific fixed effects systematically correlate with consumers' demographics. However, these results also show that price discrimination based on the demographics I observe in the microdata cannot rationalize the observed price dispersion alone. Coarse demographics only account for a

dependent variable: individual FE $\hat{\eta_i}$	(1)	(2)	(3)
Female	$902.44^{***} \\ (217.95)$		871.47^{***} (219.82)
Hispanic		534.05^{**} (221.57)	$485.16^{**} \\ (221.12)$
Mean Income _i ($$10,000s$)	$0.24 \\ (34.62)$	$12.30 \\ (36.01)$	20.02 (35.88)
Mean Purchase $\operatorname{Distance}_i$	-15.71 (9.78)	-17.45^{*} (9.85)	-15.15 (9.76)
$\begin{array}{l} \mathrm{MSA} \ \mathrm{FE} \\ \mathrm{N} \\ \mathrm{Within} \ \mathrm{R}^2 \\ \mathrm{R}^2 \end{array}$	✓ 1,150 0.015 0.082	✓ 1,150 0.009 0.076	✓ 1,150 0.020 0.087

Table 4: Demographics explain only a small fraction of price dispersion

Standard errors in parentheses

* p < 0.10, ** p < 0.05, *** p < 0.01

vanishingly small fraction of variation in the individual-specific fixed effects. More than 90% of the variation attributable to consumer fixed effects stems from unobservable heterogeneity.²⁰ These are factors unobservable to the econometrician but observable to the dealer. Dealers engage in lengthy and costly consumer interactions to extract information about their willingness to pay. The information acquisition of dealers can take the form of offering test drives, offering dealer financing to learn income and credit scores, casual chats about what the truck will be used for, or conclusions based on the general appearance or behavior of the consumer during interactions. This soft information dealers learn about consumers explains the overwhelming fraction of price dispersion, and dealers leverage it to price discriminate across consumers.

The results in table 4 also suggest that women and Hispanic customers pay significantly more for the same pickup truck, even when fully controlling for pickup truck heterogeneity. While it may be tempting to interpret these coefficients causally, note that, given the limitations of my data, I do not make this claim although these results are in line with recent research (e.g. Chandra, Gulati and Sallee, 2017).²¹ In particular, my analysis focuses on how well demographics predict the variance of the consumer-specific fixed effects. The repeat purchase data consists of observations on consumers who repeatedly buy pickup trucks and these might be non-trivially selected. For the purpose of this paper, the selection issue is mute because I only use the repeat purchase data to argue that the same consumers pay *persistently* high (low) prices across pickup truck purchases and do not focus on discrimination on gender or ethnicity. The selection based on multiple purchases is therefore irrelevant.

To conclude, the empirical facts I presented in this section highlight the importance of the use of very detailed consumer information in firms' pricing decisions for observed price dispersion. In particular, the results suggest that firms use granular consumer

 $^{^{20}{\}rm This}$ includes the contribution of MSA fixed effects. Demographics without MSA fixed effects only explain around 2% of the variation in the individual fixed effects.

²¹There exists extensive work investigating discrimination in the automobile market, see, e.g., Ayres and Siegelman (1995), Chandra, Gulati and Sallee (2017) for such studies.

information that goes beyond readily observable demographics to make personalized price offers to consumers. Whether the use of detailed consumer information benefits firms or consumers, however, is theoretically ambiguous. To understand the role of consumer information firms use for pricing in the welfare effects of price discrimination, I need a structural model.

The patterns in the data I presented motivate a structural model in which firms set personalized prices based on observable and unobservable consumer characteristics. I present such a model in the next section.

4 Model

I develop an equilibrium model of supply and demand with personalized prices to study the pickup truck market. The model shares many similarities with the class of differentiated product demand models commonly used in industrial organization following Berry, Levinsohn and Pakes (1995). However, in the model, firms set *personalized prices* based on *observed* and *unobserved* consumer heterogeneity. Thus, the model matches the main features of the pickup truck industry and is consistent with the empirical facts I presented in section 3 of this paper.

4.1 Demand

A market is defined as a year, indexed by t, and is populated by N_t consumers, each indexed by *i*. Every year, each consumer in the market chooses whether to buy a single pickup truck model *j* among the J_t available alternatives or opt for the outside option of not buying a pickup truck. I denote the outside option with j = 0 in each market *t*. Each consumer has full information on all products available and the corresponding prices she faces at the time of her choice.

I assume that each consumer maximizes her indirect utility by choosing either one

of the J_t inside goods or opting for the outside good. If consumer *i* opts for pickup truck model *j* in market *t*, her indirect utility is given by

$$u_{ijt} = \delta_{jt} + \mu_{ijt} + \varepsilon_{ijt}, \tag{2}$$

where δ_{jt} is the mean utility from choosing pickup truck model j, and μ_{ijt} the individual specific deviations from the product mean utility, both in the spirit of Berry, Levinsohn and Pakes (1995). As standard in literature, I assume that ε_{ijt} is an idiosyncratic i.i.d. distributed type I extreme value taste shock and normalize a consumer's indirect utility from choosing the outside option to $u_{i0t} = \varepsilon_{i0t}$.

While I do not specify the exact functional form of the mean utility of pickup truck j in market t, I implicitly assume that it is a function of observable pickup truck characteristics \boldsymbol{x}_{jt} and unobservable pickup truck characteristics ξ_{jt} . I depart from the standard approach in specifying the individual-specific deviations from product mean utilities μ_{ijt} . In particular, I model consumer heterogeneity as

$$\mu_{ijt} = -\alpha_{it}\rho_{ijt} + \sum_{k} \theta^{k} \boldsymbol{x}_{jt}^{k} \boldsymbol{z}_{it}^{k}.$$
(3)

Consumer heterogeneity in preferences for a pickup truck model is captured by the individual specific price sensitivity α_{it} , the individual specific price for a pickup truck ρ_{ijt} , as well as interactions between observed pickup truck characteristics \boldsymbol{x}_{jt} with observed consumer demographics \boldsymbol{z}_{it} . As is well known, these interactions capture differences in preferences for pickup truck characteristics across consumer types, allowing for substitution patterns to vary with demographics. I further assume that the support of α is the extended positive real line excluding zero.

Next, let y_{ijt} denote the choice of consumer *i*. It equals 1 if the consumer chooses pickup truck model *j* in market *t*, and 0 otherwise. Integrating over the i.i.d. type I extreme value shocks yields the familiar logit choice probabilities for each consumer *i*

$$s_{ijt} = \Pr(y_{ijt} = 1 | \boldsymbol{\rho}_{it}, \boldsymbol{x}_t, \boldsymbol{z}_{it}; \alpha_{it}, \boldsymbol{\delta}, \boldsymbol{\theta}) = = \frac{\exp(\delta_{jt} - \alpha_{it}\rho_{ijt} + \sum_k \theta^k x_{jt}^k z_{it}^k)}{\sum_{l \in J_t} \exp(\delta_{lt} - \alpha_{it}\rho_{ilt} + \sum_k \theta^k x_{lt}^k z_{it}^k)},$$

$$(4)$$

where ρ_{it} is the vector of all individual specific prices. Analogous to Berry, Levinsohn and Pakes (1995), aggregating over the individual specific choice probabilities yields aggregate choice probabilities for a given year

$$s_{jt} = \int s_{ijt} dF_{\alpha,Z}(a,z).$$
(5)

4.2 Supply

For the supply side, I assume that dealers and manufacturers act as integrated firms with aligned incentives.²² In the following, I will use dealer and firm interchangeably. I model supply as a static, full-information Nash-Bertrand pricing game between multi-product firms in which firms post personalized prices based on individual preferences for all consumers in the market. Firms can observe key preferences of consumers up to an i.i.d. type I extreme value prediction error.²³ I assume that firms are aware that they cannot learn the willingness to pay of consumers perfectly but know the distribution of the prediction errors.

I denote the set of all products within a market with \mathcal{G}_t with cardinality J_t . Let \mathcal{G}_t^d be the set of all products sold by firm d in market t. Firm d's expected profit from a single consumer i, integrating over the distribution of the i.i.d. T1EV errors, is

²²This is necessary because of data limitations. I do not have enough window sticker data for all models and dealerships in Texas to estimate a vertical model.

 $^{^{23}}$ Similar assumptions on the observability of consumer preferences for firms have recently been used in the personalized pricing literature, e.g., Buchholz et al. (2022).

given by

$$E[\pi_{i,t}^d] = \sum_{j \in \mathcal{G}_t^d} (\rho_{ijt} - c_{jt}) \cdot s_{ijt}(\boldsymbol{\rho}_{it}, \cdot).$$
(6)

where c_{jt} denotes the firm's constant marginal cost for producing and selling pickup truck model j in year t. Maximizing (6) with respect to the prices of the pickup truck models a firm offers yields the following first-order conditions

$$s_{ijt} + \sum_{j \in \mathcal{G}_t^d} (\rho_{ijt} - c_{jt}) \frac{\partial s_{ijt}}{\partial \rho_{ijt}} = 0.$$
(7)

Then, stacking all first-order conditions of all firms, one gets the usual, albeit consumer-specific, result for markups

$$\boldsymbol{\rho}_{it} - \mathbf{c}_t = (-\Omega_t \times \frac{\partial \boldsymbol{s}_{it}}{\partial \boldsymbol{\rho}_{it}})^{-1} \boldsymbol{s}_{it}, \tag{8}$$

where Ω_t represents the ownership matrix, $\frac{\partial s_{it}}{\partial \rho_{it}}$ is the matrix of all first partial derivatives of choice probabilities with respect to personalized prices, and \times represents the element-by-element multiplication.

Note that each consumer represents an independent market in this model. Because both the random coefficient α_{it} as well as demographics do not vary within consumer, each consumer has multinomial logit, not mixed logit, demand, δ_{it} . Therefore, I can use existing results on the existence and uniqueness of pure strategy Nash equilibria in literature. In particular, by theorem 7 of Konovalov and Sándor (2010), there always exists a unique Nash-equilibrium satisfying the first-order conditions (8) in $[c_t, +\infty)$.²⁴ Note that while individual-level demand functions are multinomial logit, aggregate demand is still mixed logit because price sensitivities vary across consumers.

 $^{^{24}}$ Note that Konovalov and Sándor (2010) extend the well-known results of Caplin and Nalebuff (1991) to multiproduct firms for logit demand models.

5 Identification

While the model shares many similarities with the dominant Berry, Levinsohn and Pakes (1995) approach, identification of the model's parameters differs significantly because of the individual pricing nature although some of the identification arguments still apply.

Marginal costs are identified by the minimum prices paid observed in the microdata. To understand how minimum prices for each pickup model in the data identify marginal cost, consider the following. The presence of an outside option guarantees that, as price sensitivity α approaches infinity, the prices offered converge to marginal costs. By assumption, the distribution of price sensitivities α has infinite support on the (extended) positive real line, $\alpha \in \mathbb{R}^+ \cup \{+\infty\}$. Because the logit choice probabilities ensure positive choice probabilities for any price sensitivity α , as we sample consumers, we are guaranteed to observe more and more extreme draws. In particular, for each product, the minimum price p_{jt}^{\min} observed in the data is the lowest order statistic of an i.i.d. sample of size S_{jt} , where S_{jt} is the number of micro observations for product j in market t. Therefore, the vector of minimum prices p_t^{\min} converges to the true marginal cost and thus variation in minimum prices identifies the vector of marginal cost parameters c of the model.

Identification of product mean utilities $\boldsymbol{\delta}$ and demographic preference heterogeneity parameters $\boldsymbol{\theta}$ follows the standard argument in the discrete choice demand literature (e.g. Berry, Levinsohn and Pakes, 1995). Variation in observed market shares of pickup truck models identifies $\boldsymbol{\delta}$. The parameters governing demographic heterogeneity in tastes for vehicle characteristics, $\boldsymbol{\theta}$, are identified from correlations of demographics and product characteristics with the choice of pickup truck models, which are directly observable in the microdata.

Lastly, identification of the vector of price sensitivities $\boldsymbol{\alpha}$ stems from variation

in personalized prices observable in the microdata. Conditional on buying a pickup truck model and a consumer's demographics, the residual variation in price identifies the individual-specific price sensitivity α . Note that the typical endogeneity concern in the discrete choice demand literature remains valid: personalized prices correlate with unobservable product characteristics $\boldsymbol{\xi}$. However, because prices are personalized and vary even *conditional* on pickup truck model, the inclusion of fixed effects $\boldsymbol{\delta}$ fully controls for this correlation. Therefore, contrary to the dominant Berry, Levinsohn and Pakes (1995) approach, no price instrument is needed to recover $\boldsymbol{\alpha}$.

6 Estimation

The main innovation of this paper is the framework for estimating this model. Estimation is complicated for two reasons. First, I need to jointly estimate demand and supply because prices are individual consumer-specific, and only the transaction price can ever be observed. In particular, for consumers who bought a pickup truck, the prices for the $J_t - 1$ other pickup trucks available in the market are unobservable. Second, price offers for consumers who did not buy a pickup truck can never be observed. This additionally creates a selection problem.

6.1 Estimation Strategy

To clearly distinguish between observed and unobserved prices, I define the observed, scalar-valued transaction price p_{it} as

$$p_{it} = \boldsymbol{\rho}'_{it} \cdot \boldsymbol{y}_{it},\tag{9}$$

where $\boldsymbol{\rho}_{it}$ is the $J_t \times 1$ vector of prices the individual faces in the market, and $\boldsymbol{y}_{it} = [y_{i1t}, \dots, y_{iJ_tt}]$ is a $J_t \times 1$ vector indicating the chosen pickup truck model of individual *i*: its elements equal 1 at the index of the chosen pickup truck and 0 for all other elements.

For each of the S_t consumers in market t who purchased a pickup truck model, I observe the triple $(p_{it}, \boldsymbol{y}_{it}, \boldsymbol{z}_{it})$: the transaction price, the vector indicating the consumer's choice, as well as the consumer's demographics. For each pickup truck model, I observe the tuple $(\boldsymbol{x}_{jt}, s_{jt})$: the vector of the pickup truck models' characteristics and the market share of the pickup truck model. Lastly, I denote the vector of model parameters as $\boldsymbol{\Theta} = (\boldsymbol{\theta}, \boldsymbol{\delta}, \boldsymbol{c})$. Then, I can write the model's loglikelihood as

$$\sum_{t=1}^{T} \sum_{i=1}^{S_t} \sum_{j=0}^{J_t} \mathbb{1}\{y_{ijt} = 1\} \log s_{ijt}^*(p_{it}, \boldsymbol{z}_{it}, \boldsymbol{x}_t; \boldsymbol{\Theta}) + \sum_{t=1}^{T} \sum_{j=0}^{J_t} \left(N_t s_{jt} - \sum_{i=1}^{S_t} \mathbb{1}\{y_{ijt} = 1\}\right) \log s_{jt}^*(\boldsymbol{x}_t; \boldsymbol{\Theta})$$
(10)

where s_{ijt}^* are the individual specific equilibrium choice probabilities of consumers in the microdata, and s_{jt}^* the model's predicted equilibrium market shares.²⁵ The main challenge for estimation is to recover both types of equilibrium choice probabilities to be able to form the loglikelihood. For expositional clarity, I present the strategy to form the loglikelihood for a single market, but no such restriction is necessary, and forming the loglikelihood proceeds analogously for all markets. Forming the loglikelihood given a guess of Θ proceeds roughly in 4 steps, which I describe in the following.

6.1.1 Step 1: estimating price sensitivities from microdata

To recover the unobserved prices and unobserved consumer heterogeneity α_{it} for every consumer in the microdata, I rely on the supply-side optimality conditions. I impose that the observed transaction price p_{it} satisfies the equilibrium conditions. As shown in the previous section, there exists a unique price vector satisfying the optimality

 $^{^{25}}$ A similar loglikelihood-based approach to estimate demand was recently developed by Grieco et al. (2023). Note that while the loglikelhoods share many similarities, the models are fundamentally different.

conditions

$$\boldsymbol{\rho}_{it} - \mathbf{c}_t = (-\Omega_t \times \frac{\partial \boldsymbol{\delta}_{it}}{\partial \boldsymbol{\rho}_{it}})^{-1} \boldsymbol{\delta}_{it}.$$
(11)

These J_t optimality conditions define the price vector $\boldsymbol{\rho}_{it}$ and unobserved consumer heterogeneity α_{it} as implicit functions of the model's parameters $\boldsymbol{\Theta}$ and the data. Note that since we observe one element of the price vector, this is a system of J_t equations with J_t unknowns for each consumer. Therefore, we can write α_{it} and $\boldsymbol{\rho}_{it}$ as implicit functions of observables and the models parameters

$$\alpha_{it} = \alpha(p_{it}, \boldsymbol{z}_{it}, \boldsymbol{x}_t; \boldsymbol{\Theta}),$$

$$\boldsymbol{\rho}_{it} = \boldsymbol{\rho}^p(p_{it}, \boldsymbol{z}_{it}, \boldsymbol{x}_t; \boldsymbol{\Theta}).$$
 (12)

Next, because these functions are defined implicitly only, we have to numerically solve the first-order conditions of firms in (8) for the corresponding vector of prices ρ_{it} and price sensitivity α_{it} . This requires solving J_t equations for J_t unknowns.

However, instead of directly using (8) to solve for J_t prices and price sensitivity, we can rely on a well-known property of multinomial logit demand models: firms set the same profit margin, defined as price minus marginal cost, for all products. Therefore, we only need to solve for D_t instead of J_t unknowns, where D_t is the number of firms active in market t. This considerably speeds up estimation. After solving for α_{it} for all consumers, we have the vector of consumers' price sensitivities conditional on purchase.

Substituting the functions in (12) for prices and price sensitivity also allows us to rewrite the choice probabilities s_{ijt} for consumers in the microdata entirely in terms of data and parameters imposing the optimality conditions of the firms

$$s_{ijt}^{*}(p_{it}, \boldsymbol{z}_{it}, \boldsymbol{x}_{t}; \boldsymbol{\Theta}) = \frac{\exp(\delta_{jt} - \alpha(p_{it}, \boldsymbol{z}_{it}, \boldsymbol{x}_{t}; \boldsymbol{\Theta}) \cdot \rho_{j}^{p}(p_{it}, \boldsymbol{z}_{it}, \boldsymbol{x}_{t}; \boldsymbol{\Theta}) + \sum_{k} \theta^{k} x_{jt}^{k} z_{it}^{k})}{1 + \sum_{l \in J_{t}} \exp(\delta_{lt} - \alpha(p_{it}, \boldsymbol{z}_{it}, \boldsymbol{x}_{t}; \boldsymbol{\Theta}) \cdot \rho_{k}^{p}(p_{it}, \boldsymbol{z}_{it}, \boldsymbol{x}_{t}; \boldsymbol{\Theta}) + \sum_{k} \theta^{k} x_{lt}^{k} z_{it}^{k})}$$
(13)

Thus, after this step, we have the distribution of price sensitivities conditional on purchase and demographics, as well as the vector of choice probabilities for each consumer in the microdata. In the next step, we need to solve the selection into purchase problem.

6.1.2 Step 2: estimating the distribution of price sensitivities conditional on demographics

Let \mathcal{Y} the random variable indicating whether a consumer bought any pickup truck model or opted for the outside good. Its realization y equals 1 if the consumer chooses a pickup truck model, and 0 if she chooses the outside option.

From step 1, we have the vector of consumers' price sensitivities conditional on purchase and demographics. I first estimate the distribution of price sensitivities conditional on purchase and demographics, $dF_{\alpha|y=y,Z=z}(a|y \neq 0, Z = z)$.²⁶ To form predicted market shares, however, we need the distribution of price sensitivities conditional on demographics only.

To solve the selection into purchase problem, we can again exploit the optimality conditions in (8): for every point (a, \mathbf{z}) in the support of the joint distribution of price sensitivities and demographics, $F_{\alpha,Z}(a, \mathbf{z})$, the supply side optimality conditions uniquely pin down the corresponding price vector

$$\boldsymbol{\rho} = \boldsymbol{\rho}^{\alpha}(a, \boldsymbol{z}, \boldsymbol{x}; \boldsymbol{\Theta}). \tag{14}$$

²⁶Note that here I assume that demographics are following a discrete distribution. While none of this is necessary and can be relaxed, my demographic data comes in bins.

As we will see, this relation is key to recovering the distribution of price sensitivities conditional on demographics only. From the functional form of the logit choice probabilities, it follows that

$$P(y=0|\alpha=a, Z=\boldsymbol{z}) = \frac{1}{1 + \sum_{l=1}^{J_t} \exp(\delta_{lt} - a \cdot \rho_j^{\alpha}(a, \boldsymbol{z}, \boldsymbol{x}_t; \boldsymbol{\Theta}) + \sum_k \theta^k x_{lt}^k z^k)}, \quad (15)$$

where ρ_l^{α} is the *l*-th element of the implicit pricing function from the supply side first-order conditions defined in (14).

Next, $Pr(y \neq 0 | Z = z)$ is the sum of the conditional inside shares for a given demographic group. We can express these as a function of the data by an application of Bayes' rule

$$P(y \neq 0 | Z = \boldsymbol{z}) = \frac{P(\boldsymbol{z} | y \neq 0) P(y \neq 0)}{P(Z = \boldsymbol{z})},$$
(16)

where $P(\boldsymbol{z}|\boldsymbol{y}\neq 0)$ can be estimated directly from the microdata, $P(\boldsymbol{y}\neq 0)$ are the unconditional inside shares from the transaction level data, and $P(\boldsymbol{Z}=\boldsymbol{z})$ is the joint distribution of demographics in Texas from the Current Population Survey.

By applying Bayes' rule, we can now obtain the distribution of price sensitivities conditional on demographics only, solving the selection into purchase problem

$$dF_{\alpha|Z=z}(a|Z=z) = \frac{dF_{\alpha|Y=y,Z=z}(a|y\neq 0, Z=z)P(y\neq 0|Z=z)}{1-P(y=0|\alpha=a, Z=z)}.$$
 (17)

Now, we can use this distribution to form the model's predicted shares in the next step.

6.1.3 Step 3: forming model's predicted shares

To obtain predicted market shares, we need to integrate over the individual-specific choice probabilities of consumers in the market. Using (14), the choice probabilities at the point (a, z) are given by

$$s_{jt}^{a,z}(a, \boldsymbol{z}, \boldsymbol{x}_t; \boldsymbol{\Theta}) = \frac{\exp(\delta_{jt} - a \cdot \rho_j^{\alpha}(\boldsymbol{z}, \boldsymbol{x}_t, a; \boldsymbol{\Theta}) + \sum_k \theta^k x_{jt}^k z^k)}{1 + \sum_{l \in J_t} \exp(\delta_{lt} - a \cdot \rho_l^{\alpha}(\boldsymbol{z}, \boldsymbol{x}_t, a; \boldsymbol{\Theta}) + \sum_k \theta x_{lt}^k z^k)}.$$
 (18)

Then, integrating over the joint distribution of price sensitivities and consumer demographics yields the predicted market shares $s_{jt}^*(\cdot)$:

$$s_{jt}^{*}(\boldsymbol{x}_{t};\Theta) = \int s_{jt}^{a,z}(a,\boldsymbol{z},\boldsymbol{x}_{t};\Theta) dF_{\alpha,Z}(a,\boldsymbol{z})$$

$$= \int s_{jt}^{a,z}(a,\boldsymbol{z},\boldsymbol{x}_{t};\Theta) dF_{\alpha|Z=\boldsymbol{z}}(a|Z=\boldsymbol{z}) dF_{Z}(\boldsymbol{z}),$$
(19)

which are the model analogs to observed market shares in the data. Here, we used the distribution of price sensitivities conditional on demographics from equation (17).

6.1.4 Step 4: forming and maximizing the Loglikelihood

From step 1, we have the choice probabilities s_{ijt}^* of the consumers in the microdata; from step 3 we have the model's predicted market shares s_{jt}^* for all products. Thus, we can form the model's loglikelihood for market t by

$$\log \mathcal{L}_{t}(\boldsymbol{p}_{t}, \boldsymbol{z}_{t}, \boldsymbol{x}_{t}, \boldsymbol{s}_{t}; \boldsymbol{\Theta}) = \sum_{i=1}^{S_{t}} \sum_{j=0}^{J_{t}} \mathbb{1}\{y_{ijt} = 1\} \log s_{ijt}^{*}(p_{it}, \boldsymbol{z}_{it}, \boldsymbol{x}_{t}; \boldsymbol{\Theta}) + \sum_{j=0}^{J_{t}} \left(N_{t}s_{jt} - \sum_{i=1}^{S_{t}} \mathbb{1}\{y_{ijt} = 1\}\right) \log s_{jt}^{*}(\boldsymbol{x}_{t}; \boldsymbol{\Theta}).$$
(20)

Note that I did not impose that the distribution of price sensitivities is the same

across markets.²⁷ The model's loglikelihood is then the sum of all the markets' loglikelihoods,

$$\log \mathcal{L}(\cdot; \boldsymbol{\Theta}) = \sum_{t=1}^{T} \log \mathcal{L}_t(\cdot; \boldsymbol{\Theta}).$$
(21)

In practice, not every parameter guess will be consistent with all first-order conditions of all consumers in my data. Therefore, if a parameter guess creates such a problem, I add a smooth, quadratic penalty term to the loglikelihood.

Note that it is possible to reduce the dimension of the parameter vector the optimization routine has to search over. In particular, one can reduce the computational burden by using a two-step estimator. In the first step, one can estimate marginal cost \boldsymbol{c} directly from the data using the minimum observed prices for each pickup truck model as cost estimates because minimum prices $\boldsymbol{p}_{jt}^{\min}$ for each product j and market t converge to the true marginal cost. In the second step, one then uses the loglikelihood to estimate the remaining parameters in $\tilde{\Theta} = \{\boldsymbol{\theta}, \boldsymbol{\delta}\}$. This reduces the parameters to be estimated in the search over the loglikelihood by $J = \sum_m J_m$. I use this version of my estimator for estimation.

6.1.5 Measuring the correlation of price sensitivity and consumer demographics

After following the procedure outlined above, we have estimates for Θ and $F_{\alpha,Z}(a, z)$. Note that these can now be used to explore the correlation between observed consumer demographics and price sensitivities in a second estimation step. Since I did not assume that α_{it} follows a specific parametric distribution, we can sample from $\hat{F}_{\alpha,Z}(a, z)$ to understand how price sensitivities and demographics correlate. Given the estimates

²⁷While this allows me to consistently estimate the distribution of α if the distribution of α is the same across markets, a more efficient estimator that does not separate estimation by market in the inner loop would exist. Hence, my estimator is only efficient if the distribution of price sensitivities differs across years.

	pickup truck characteristics					
	constant	Size	Miles/Gal.	US-Brand		
income	0.498 (0.025)	$0.145 \\ (0.017)$	$0.090 \\ (0.018)$	-0.214 (0.030)		
female	-1.962 (0.020)					
Hispanic	-0.719 (0.020)	_		_		

Table 5: Parameter Estimates of Structural Model

Note: Standard errors in parentheses.

All continuous pickup truck characteristics and income standardized.

for Θ , we can draw from the estimated joint distribution. After that, one can then regress these draws on the demographics to understand how price sensitivities and consumer demographics correlate or plot the corresponding distributions.

7 Results

I estimate my model using the *microdata for demand estimation* dataset covering 2016 – 2019 I described in section 2. Note that the *microdata for demand estimation* is normalized to the model level as described in section 2.

7.1 Empirical Specification

I specify consumer *i*'s indirect utility from purchasing pickup truck model j in market t as

$$u_{ijt} = \delta_{jt} - \alpha_{it}\rho_{ijt} + \theta_1 \text{income}_{it} + \theta_2 \text{size}_{jt} \times \text{income}_{it} + \theta_3 \text{mpg}_{jt} \times \text{income}_{it} + \theta_4 \text{US-Brand}_{jt} \times \text{income}_{it} + \theta_5 \text{female}_{it} + \theta_6 \text{Hispanic}_{it} + \varepsilon_{ijt}$$
(22)

where income is standardized to have zero mean and unit variance, size is the length times the height of a pickup truck, mpg is the miles per gallon rating, and US-Brand is an indicator variable for US makes. I standardize all continuous pickup truck characteristics. This specification allows consumers of different sociodemographic groups to differ in their taste for buying a pickup truck model. I do not parametrize the distribution of α_{it} but nonparametrically estimate it from the microdata exploiting the supply side equilibrium conditions in (8). Price sensitivities may be arbitrarily correlated with consumer demographics. The advantages of this flexible specification are two-fold. First, it allows the model to nonparametrically recover the consumer information firms use in their pricing decisions. Second, it allows for asymmetries in firms' rankings of consumers with respect to their demand elasticities in the spirit of **Corts** (1998).²⁸

7.2 Parameter Estimates

Table 5 presents the results for the estimates of the structural parameters θ . These parameters govern the taste for product characteristics by consumer demographics. The results intuitively match the correlations in the microdata uncovered in section 3. Higher-income individuals are likelier to purchase a pickup truck in any given year than low-income individuals. Female buyers are less likely to buy a pickup truck than

 $^{^{28}}$ This corresponds to what Corts (1998) coined the best response asymmetry.

males. Similarly, Hispanic consumers are less likely to purchase a pickup truck in any given year compared to non-Hispanic consumers.

7.3 Estimates of Price Sensitivities and Correlations with Demographics

The estimates for α_{it} show consumers' strong distaste for price. I estimate the mean price sensitivity to be 2.95, with a standard deviation of 2.11, highlighting the substantial unobserved heterogeneity in price sensitivity. Figure 3 plots the estimated *unconditional* distribution of price sensitivities in 2018 illustrating the substantial dispersion of price sensitivities.²⁹



Figure 3: Unconditional distribution of price sensitivities

Contrary to figure 3, figure 4 plots the *conditional* distributions of price sensitivities for selected demographics (females, males, Hispanics, and non-Hispanics) in 2018.³⁰

²⁹Note that I do not impose that the distribution of price sensitivities is constant across years. Figure 9 for 2017 in the appendix shows that there is variation across years, but the results are qualitatively similar to 2018.

 $^{^{30}}$ As for the unconditional distribution, figure 10 in the appendix plots the conditional distribution of price sensitivities in 2017 for reference.



Figure 4: Conditional distributions of price sensitivities

These plots highlight the substantial heterogeneity in price sensitivity even conditional on consumers' demographics. The mean differences in price sensitivities across these groups intuitively match the findings in section 3. While the mean of the distribution of price sensitivities of both females and Hispanics is lower than that for males and non-Hispanics respectively, the distributions of price sensitivities considerably overlap. This implies that while mean differences across demographic groups exist, heterogeneity within demographic groups is much more significant.

7.4 Marginal Costs

The average estimated marginal cost for a pickup truck model is \$18,911. This aligns well with industry reports. The estimates exhibit intuitive rankings regarding manufacturers: Toyota produces the most expensive pickup trucks, while Nissan produces the cheapest. These estimates match the price levels observed in the transaction level data well.

I project estimated marginal costs of producing a pickup truck on pickup truck

	marginal cost			
Size	0.495^{***} (0.167)			
MPG	0.181^{***} (0.064)			
Horsepower	$0.001 \\ (0.001)$			
US-Brand	-1.185^{***} (0.269)			
R^2 Year FE	0.501 ✓			
Note: Standard errors in parentheses				

* p < 0.10, ** p < 0.05, *** p < 0.01Size in thousands of square inches.

Table 6: Regression of marginal costs on pickup truck characteristics

characteristics. Results are reported in table 6. Pickup trucks of American manufacturers are estimated to be cheaper to produce than foreign pickup trucks. Additionally, larger pickup trucks are more costly to produce. Trucks with higher per gallon ratings are more expensive to produce but horsepower is estimated to have no significant effect on marginal cost.

7.5 Price Elasticities and Markups

A complication for calculating price elasticities and markups given the transactionby-transaction pricing decisions of firms is that separate own-price elasticities exist for each consumer, which is non-standard. To obtain a comparable statistic to the standard notion of an own-price elasticity, I define the market-level own-price elasticity for pickup truck model j in market t as follows

$$\epsilon_{jt} = -\int a \cdot \rho_j^{\alpha}(\cdot)(1 - s_{jt}^{a,z}(\cdot))dF_{\alpha,Z}(a,z).$$
(23)

The interpretation of this statistic changes relative to the standard notion of price elasticity in the following way: It measures the change in demand for product j when the average price level in the market for product j increases by 1%. As defined above, the estimated parameters imply that the average market-level own-price elasticity across all products and markets is -7.31. Therefore, demand is relatively more elastic than in previous studies on the automobile market. However, this model's more priceelastic demand curve is not surprising since preceding studies have relied on MSRPs or average transaction prices as proxies for all quoted prices in a market. As it turns out, the consumers who ultimately purchase a pickup truck are the least price-sensitive consumers in the market, paying the highest markups. Using the average transaction price thus significantly overstates the quoted prices in the market. Consequently, these studies' estimated mean price sensitivity parameters can be severely downward biased. Similarly, as shown in section 3, most purchasers pay significantly less than MSRP for their pickup trucks. Therefore, using MSRP exacerbates the problems that the average transaction price induces if the MSRP is larger than the average transaction price.

Similar to the definition of the market-level own-price elasticity for a pickup truck model, I define the market-level markup for pickup truck model j in market t as

$$\operatorname{markup}_{jt} = \int \frac{\rho_j^{\alpha}(\cdot) - c_{jt}}{\rho_j^{\alpha}(\cdot)} dF_{\alpha,Z}(a,z).$$
(24)

This definition of the average market-level markup for a pickup truck model is analogous to the familiar product-level markups measured by the Lerner index. I estimate that the average markup is around 0.21, while the average share-weighted markup is 0.35. The difference in both again highlights the importance of distinguishing between the models' estimated price quotes for the whole market and the realized prices in the transaction data. Since most consumers in the market are very price-sensitive, average markups are substantially lower than average share-weighted markups.

Lastly, to compare the estimates of the model with industry sources on price-cost margins, I define the average price-cost margin for pickup truck model j in market t as

$$\operatorname{margin}_{jt} = \int (\rho_j^{\alpha}(\cdot) - c_{jt}) dF_{\alpha, Z}(a, z).$$
(25)

Again, price-cost margins are consumer-specific, but we can define a market-level analog to the familiar statistic. Using the definition of the market-level price-cost margin for pickup truck model j, I estimate that the average market-level margin across all products is around \$5,592. The average share-weighted margin is \$10,683. These margins might seem high for the automobile market in general, but it is important to remember that I study only pickup trucks. In particular, my estimates are in line with industry reports. For example, industry reports suggest that Ford makes around \$13,000 in profit per pickup truck (Automotive News, 2015).

To sum up, the model produces reasonable estimates for demand, costs, and markups. While I have to resort to new definitions for elasticities and margins, these statistics align well with what has been reported in industry sources. In the next section, I will use these estimates to study the welfare effects of personalized pricing.

8 The Value of Consumer Information for Discriminatory Pricing

Are personalized prices benefitting or harming firms and/or consumers? A crucial component in answering this question is the amount of consumer information firms can use to personalize prices and price discriminate. To isolate the role of information in the welfare effects of price discrimination, I conduct multiple counterfactual simulations. Each of the counterfactuals uses a specific form of pricing in which firms can use *less* consumer information than the actual amount of consumer information firms use to personalize prices under the baseline. In particular, I consider the following counterfactual simulations:

1. Uniform pricing:

First, I conduct a counterfactual in which firms cannot use any information beyond the distributions of price sensitivities and demographics in the population. Firms then set a single price for all consumers for each pickup truck model they produce and sell, as in the canonical Berry, Levinsohn and Pakes (1995) model. It is well known that under this informational paradigm, the stacked first-order conditions take the form of

$$\boldsymbol{\rho}_t^u - \mathbf{c}_t = \left(-\boldsymbol{\Omega}_t \times \frac{\partial \boldsymbol{s}_t}{\partial \boldsymbol{\rho}_t^u}\right)^{-1} \boldsymbol{s}_t, \tag{26}$$

where $\boldsymbol{\rho}_t^u$ is the $J_t \times 1$ vector of uniform prices. These first-order conditions define the vector of uniform prices $\boldsymbol{\rho}_t^u$ as an implicit function of the model's parameters and the data.³¹ I obtain the counterfactual price vector $\boldsymbol{\rho}_t^u$ by numerically solving the first-order conditions.

2. Price discrimination on gender, ethnicity, and income only:

Second, I conduct a counterfactual in which firms have information on gender, ethnicity, and income of a consumer, but firms do not observe the consumerspecific price sensitivity. However, they know the distribution of price sensitivities even conditional on consumer demographics. Firms then set a single price for each pickup truck they produce and sell for each demographic group. For example, firms set a single price for all Hispanic females with incomes between \$50,000 and \$75,000 per year for each pickup truck model they offer. This corresponds to the baseline of the dominant approach studying price discrimination in the automobile market (e.g. Langer, 2016; D'Haultfoeuille, Durrmeyer and Février, 2019). The corresponding stacked first-order conditions are equivalent to (26) but are now *demographic group-specific*

$$\boldsymbol{\rho}_t^d - \mathbf{c}_t = (-\boldsymbol{\Omega}_t \times \frac{\partial \boldsymbol{s}_t^d}{\partial \boldsymbol{\rho}_t^d})^{-1} \boldsymbol{s}_t^d, \qquad (27)$$

where $\boldsymbol{\rho}_t^d$ is the $J_t \times 1$ vector of discriminatory prices for demographic group d, and \boldsymbol{s}_t^d the model's predicted aggregate choice probabilities for demographic group d.

 $^{^{31}}$ It is well known that a Nash equilibrium does not need to exist in random coefficients discrete choice models. Therefore, I assume the existence of a pure strategy Nash equilibrium satisfying the first-order conditions as is standard in literature (e.g. Berry, Levinsohn and Pakes, 1995).

To quantify the consumer welfare changes under the different informational paradigms, I use the compensating variation. In particular, the compensating variation in market t comparing consumer surplus under uniform prices to personalized pricing is given by

$$\Delta CS_t = \int \frac{1}{a} \log(1 + \sum_{j \in \mathcal{G}_t} \exp(\delta_{jt} - a\rho_{jt}^u + \sum_k \theta^k x_{jt}^k z^k)) dF_{\alpha,Z}(a, z) - \int \frac{1}{a} \log(1 + \sum_{j \in \mathcal{G}_t} \exp(\delta_{jt} - a\rho_j^{a,z}(\cdot) + \sum_k \theta^k x_{jt}^k z^k)) dF_{\alpha,Z}(a, z).$$

$$(28)$$

As usual, I measure welfare changes for firms in profits.

8.1 Counterfactual results

The counterfactual simulations allow me evaluate the welfare effects of firms' use of consumer information beyond the distribution of demographics and price sensitivities in population. First, the uniform pricing counterfactual isolates the value of consumer information to firms relative to only knowing the distribution of consumer characteristics including price sensitivities in population. Second, the counterfactual in which firms can only use demographic consumer information observable to the econometrician and the conditional distributions of price sensitivities isolates the importance of observable versus unobservable consumer heterogeneity.

The value of consumer information for personalized pricing: The *personalized pricing* columns of table 7 show that firms benefit significantly from using information beyond the distribution of demographics and price sensitivities in population. Acquiring detailed, fine-grained information through direct interactions with individual consumers and using this information to charge consumer-specific, personalized prices substantially raises profits by around 19% when compared to profits under uniform pricing. In particular, not only aggregate industry profits but profits across all manufacturers are unequivocally higher than under uniform pricing. While the

	price discr personaliz	imination ed pricing	price di gender, ethn	price discrimination gender, ethnicity, income only		
	\$ millions	%	\$ millions	%		
Toyota	+75.67	+22.57%	-0.93	-0.37%		
Honda	+4.06	+23.37%	-0.01	-0.05%		
Nissan	+10.86	+15.73%	-0.17	-0.26%		
Ford	+72.75	+14.76%	-1.10	-0.26%		
GM	+90.65	+14.59%	-2.11	-0.40%		
RAM	+36.04	+16.69%	-0.27	-0.14%		
Industry profits	+290.05	+19.75%	-4.59	-0.31%		

Table 7: Average gains/losses in profit across years relative to uniform pricing due to price discrimination on all consumer characteristics (personalized pricing) and gender, race, and income only.



Figure 5: Changes in yearly profit per manufacturer by personalizing prices in \$ Millions (left panel) and percentages (right panel) when compared to uniform pricing.

profitability of price discrimination based on detailed consumer features is unambiguous across firms, there is however heterogeneity in how much firms profit from using this information. Figure 5 visualizes this heterogeneity across manufacturers. These results on the profitability of obtaining and using detailed consumer information for personalizing prices hint at why dealers spend considerable effort to interact with consumers in the automobile industry when this does not happen in other industries: it is highly profitable to exert this effort for comparatively expensive consumer goods.

Relative to uniform pricing, the increase in industry profits from using consumer information beyond coarse demographic groups stems from two sources. First, dealers leverage consumer information to offer price-sensitive consumers better deals than under uniform pricing. Thus, under personalized pricing, price-sensitive consumers are likelier to buy a pickup truck than under uniform prices. This realizes profits dealers could not have earned under uniform pricing. Second, dealers can use the same consumer information to charge price-insensitive consumers higher prices than under uniform prices while making them less likely to buy pickup trucks. Figure 7 illustrates how consumers' choice probabilities across different levels of price sensitivities change



Figure 6: Changes in yearly profit per manufacturer by using discriminatory pricing based only on demographics in \$ Millions (left panel) and percentages (right panel) when compared to uniform pricing.

when moving from uniform pricing to personalized pricing. These sources have opposing forces on the average prices offered in the market, as well as on aggregate demand. While personalized pricing, compared to uniform pricing, reduces prices on *average* by 15%, share-weighted prices increase by 2%. Share-weighted prices increase because price-insensitive consumers are more likely to buy pickup trucks. Finally, personalized pricing also leads to higher demand. On average, across years, dealers' use of personalized pricing increases pickup truck sales by around 10% when compared to uniform pricing.

Next, I show that while personalizing prices is highly profitable to firms, price discrimination in the pickup truck market need *not* raise profits. The *gender*, *ethnicity*, *income only* columns of table 7 show that different information structures can completely reverse the profitability of using individual consumer information for pricing if information is coarse. In particular, industry profits from price discrimination only on gender, ethnicity, and income, and using the conditional distribution of price sensitivities, would be around 0.3% lower than under uniform pricing. Figure 6 furthermore



Figure 7: Probability of purchasing any pickup truck model under personalized pricing (red) and uniform pricing (blue) in 2017 for a non-Hispanic male as a function of price sensitivity.

visualizes that when firms are restricted to using information on demographic groups only, *all* firms lose from price discrimination. Notably, these results are qualitatively in line with findings of recent literature studying price discrimination in the US automobile market based on demographic groups (Langer, 2016).³²

These results highlight that dealers' primary source of profits is soft information beyond coarse demographic groups. While dealers use information on demographic groups and income for pricing, these demographics only offer limited information about consumers' willingness to pay because the dispersion of price sensitivity is significant, even within demographic groups. This large heterogeneity in price sensitivities makes personalization of price offerings highly profitable for dealers. This helps to rationalize why sales personnel invest considerable effort into learning the willingness to pay of consumers via long conversations during test drives, the negotiations phase, and even

 $^{^{32}}$ Using a model with price discrimination based on coarse demographic groups, Langer (2016) finds that price discrimination based on gender and marital status would reduce industry profits by around 3% per year.

financial information when running consumers' credit. Additionally, these results also show how complicated the relationship between the informational paradigm and industry profits from price discrimination is. In particular, these results show that while firms profit immensely from having access to very detailed consumer information, in a competitive environment, firms can be worse off than having access to no such information at all when only slightly changing the informational environment. Such situations could, for example, arise via the introduction of legislation limiting the amount of consumer information firms can collect but still allows firms to learn some basic information on consumers' demographics, such as the European GDPR.

The result that soft information drives the profitability of personalized pricing also shows why personalized pricing did not lead to lower prices in the pickup truck market. Corts (1998) showed that competitive third-degree price discrimination can lower profits. These results rest on a property that Corts (1998) called best response asymmetry. My model is flexible enough to capture these asymmetries. For example, my estimates imply that higher-income consumers prefer larger, foreign-made pickup trucks while low-income consumers prefer smaller, domestic ones. US manufacturers have an incentive to offer lower prices to high-income consumers, while foreign firms have the same incentives for low-income consumers. Corts (1998) showed that these asymmetries can lead to more pricing pressure and lower prices for all consumers. However, the estimates also imply that the drivers of price dispersion in the pickup truck market are differences in price sensitivities across consumers. The institutional details of the market allow dealers to learn these price sensitivities as soft information and price accordingly. However, all dealers use price sensitivities similarly: less pricesensitive consumers are offered higher prices. Because price sensitivities cannot induce best response asymmetries, personalized pricing did not lead to lower prices for all consumers in the pickup truck market.



Figure 8: Losses in consumer surplus due to personalized pricing by demographic group. Left panel shows average percentage change in consumer surplus, right panel percentage of people in demographic group worse off than under uniform prices.

Effect of Personalized Pricing on Consumer Surplus: Price discrimination on very detailed levels of consumer information reduces aggregate consumer surplus. In particular, as figure 8 illustrates, price personalization by dealers decreases *average* compensating variation by around 7% relative to a situation with posted uniform prices. This corresponds to losses of around \$198 million in consumer surplus per year. Figure 8 also shows that, on average, there is substantial heterogeneity in consumer surplus losses from personalized pricing across demographic groups.

While the average consumer in the pickup truck market loses from dealers using granular consumer information to personalize prices, most consumers, however, are better off. Some consumers lose a lot, while most gain little compared to uniform prices. The left panel in figure 8 shows that the losses in consumer surplus stem from a minority of around 19% of consumers. Moreover, even though female and Hispanic consumers lose the most on average, the majority of female and Hispanic consumers still benefit from personalized pricing.

How could total consumer surplus decrease if only a small fraction of consumers are worse off and most are quoted lower prices by dealers? Due to dealers' ability to infer consumers' willingness to pay, price-insensitive consumers face much higher prices for pickup trucks under personalized pricing than under uniform prices. On the other hand, price-sensitive consumers profit from lower prices. However, the gains for very price-sensitive consumers from personalized pricing quickly become negligible with increasing values of α_{it} . This is because these consumers rarely buy a pickup truck even when facing prices that are very close to firms' marginal costs. Whether or not aggregate consumer surplus increases or decreases when dealers personalize prices ultimately depends on the precise distribution of price sensitivities. As it turns out, there are only around 19% of consumers who lose a lot from discriminatory pricing. About 30% of consumers gain little from discriminatory pricing, while the rest of consumers see only a vanishingly small, but positive, change in their consumer surplus. In aggregate, losses for the 19% of price-insensitive consumers outweigh the gains of most consumers, and thus aggregate consumer surplus decreases.

Figure 7 shows why some consumers are better off, why some are worse off, and why some consumers gain little even from very competitive prices. It plots the probability of buying any pickup truck model in a given year as a function of price sensitivity for a non-Hispanic male in 2017. The figure shows that the probability of purchasing any pickup truck model under personalized pricing (price discrimination) is rotated around a point for which consumers are indifferent between dealers engaging in personalized pricing and uniform pricing. When dealers personalize prices, priceinsensitive consumers' probability of purchasing any pickup truck model decreases. On the other hand, the probability of purchase increases for more price-sensitive consumers. Since compensating variation is a function of purchase probabilities, less price-sensitive consumers lose while more price-sensitive consumers benefit. However, with increasing price sensitivities the difference between the two pricing models quickly vanishes.

Effect of Personalized Pricing on Total Surplus: Lastly, it is important to notice that while personalized prices harm consumers *on average*, reducing aggregate consumer surplus, the losses in aggregate consumer surplus are less than the gains in profit for firms. This aligns well with basic economic intuition: since dealers can leverage very granular consumer information for their price quotes, prices are closer to the actual willingness to pay of consumers leading to a more efficient market outcome. Note that firms still face uncertainty about consumers' willingness to pay because firms cannot observe the T1EV error term and prices do not exactly equal consumers' willingness to pay as with first-degree price discrimination. Since firms' gains in profits from personalized pricing outweigh losses in aggregate consumer surplus when compared to uniform pricing, across years, the average total surplus increases by around \$92 million per year.

9 Conclusions

In this paper, I showed that dealers leverage large amounts of consumer information when making pricing decisions. Fully controlling for the heterogeneity of pickup trucks and leveraging panel data on consumers, I showed that dealers use consumer information that goes far beyond demographics to personalize prices. Almost all personalization of prices is based on this soft information rather than demographics. Using a novel equilibrium model of supply and demand with personalized pricing, I show that personalized pricing increases profits and total surplus, but *on average* harms consumers. However, whether or not access to consumer information increases profits for firms crucially depends on how detailed these data are. In particular, I show that if firms can only price discriminate on readily observable demographics, firms lose from price discrimination relative to a situation in which they have less information and charge uniform prices.

These results on how the granularity of consumer information drives welfare results complement the recent theoretical literature studying personalized pricing with evidence from the United States' second-most important consumer goods market. I show that the frequently raised concerns about restricting firms from employing personalized prices do not apply in the US automobile market. In particular, allowing firms to engage in price personalization leads to a reduction in average consumer surplus in the Texas market for pickup trucks. However, there is substantial heterogeneity in who carries the burden of consumer welfare losses. Personalization of prices redistributes surplus from price-insensitive consumers to price-sensitive consumers as well as firms. I also show that while there are winners and losers from personalized pricing in the pickup truck market, personalized prices in this market increase total welfare. Theoretical literature has long echoed concerns about the possible consumer welfare losses from restricting personalized pricing. My results provide contrasting real-world evidence but also call for more empirical studies validating these concerns in other markets, in particular in online settings.

Lastly, there are many fruitful directions for future research. My approach focuses on the main features of the pickup truck market. Although I have detailed and complete data on the consumer and product level for the US automobile market, I cannot consider markets closely tied to the market for pickup trucks. An exciting direction for future research might be combining the individualized pricing of this model with the trade-in or financing negotiations process. Prices in these adjacent markets are usually highly correlated with the actual transaction price in the new pickup truck market and might offer additional insights about consumers' price sensitivities and firms' information about consumers' willingness to pay.

References

- Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica*, 67(2): 251–333.
- Automotive News. 2015. "The F-150, Fords heavy-duty profit hauler." https: //www.autonews.com/article/20150428/BL0G06/150429797/the-f-150-for d-s-heavy-duty-profit-hauler, Last accessed on 2023-07-01.
- Ayres, Ian, and Peter Siegelman. 1995. "Race and Gender Discrimination in Bargaining for a New Car." *American Economic Review*, 3(85): 304–321.
- Bernstein, Shai, Rebecca Diamond, Abhisit Jiranaphawiboon, Timothy McQuade, and Beatriz Pousada. 2022. "The Contribution of High-Skilled Immigrants to Innovation in the United States." National Bureau of Economic Research Working Paper 30797.
- Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile prices in market equilibrium." *Econometrica*, 841–890.
- Buchholz, Nicholas, Laura Doval, Jakub Kastl, Filip Matejka, and Tobias Salz. 2022. "The Value of Time: Evidence from Auctioned Cab Rides." Princeton University. Economics Department. Working Papers 2022-22.
- Caplin, Andrew, and Barry Nalebuff. 1991. "Aggregation and Imperfect Competition: On the Existence of Equilibrium." *Econometrica*, 59(1): 25–59.
- Chandra, Ambarish, Sumeet Gulati, and James M. Sallee. 2017. "Who Loses When Prices Are Negotiated? An Analysis of the New Car Market." *The Journal* of Industrial Economics, 65(2): 235–274.

- **Corts, Kenneth.** 1998. "Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment." *The RAND Journal of Economics*, 29(2): 306–323.
- Coşar, A. Kerem, Paul L. E. Grieco, and Felix Tintelnot. 2015. "Borders, Geography, and Oligopoly: Evidence from the Wind Turbine Industry." *The Review* of Economics and Statistics, 97(3): 623–637.
- Cox Automotive. 2019. "2019 Car Buyer Journey Study." https://www.coxautoi nc.com/learning-center/2019-car-buyer-journey-study/, Last accessed on 2023-06-10.
- Crawford, Gregory S., Nicola Pavanini, and Fabiano Schivardi. 2018. "Asymmetric Information and Imperfect Competition in Lending Markets." The American Economic Review, 108(7): 1659–1701.
- D'Haultfoeuille, Xavier, Isis Durrmeyer, and Philippe Février. 2019. "Automobile Prices in Market Equilibrium with Unobserved Price Discrimination." *The Review of Economic Studies*, 86(5): 1973–1998.
- Diamond, Rebecca, Tim McQuade, and Franklin Qian. 2019. "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco." American Economic Review, 109(9): 3365–94.
- Drury, Ivan, and Jessica Caldwell. 2020. "edmunds Truck Report 2020." https: //static.ed.edmunds-media.com/unversioned/img/industry-center/insigh ts/2020-truck-report.pdf, Last accessed on 2023-08-04.
- Dubé, Jean-Pierre, and Sanjog Misra. 2023. "Personalized Pricing and Consumer Welfare." Journal of Political Economy, 131(1): 131 – 189.

- Goldberg, Pinelopi Koujianou. 1996. "Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey." Journal of Political Economy, 3(104): 622–654.
- Grieco, Paul, Charles Murry, Joris Pinkse, and Stephan Sagl. 2023. "Conformant and Efficient Estimation of Discrete Choice Demand Models." Working Paper.
- Grigolon, Laura, Liana Jacobi, and Michelle Sovinsky. 2020. "The Price is (not) Right: Incorporating Unobserved Price Heterogeneity in Demand Models." Working Paper.
- Grunewald, Andreas, Jonathan Lanning, David Low, and Tobias Salz. 2023.
 "Auto Dealer Loan Intermediation: Consumer Behavior and Competitive Effects."
 Working Paper.
- Harless, David W., and George E. Hoffer. 2002. "Do Women Pay More for New Vehicles? Evidence from Transaction Price Data." *American Economic Review*, 1(92): 270–279.
- Infutor Data Solutions. 2023. "Auto Data Appends." https://infutor.com/prof ile-appends/automotive-appends/, Last accessed on 2023-06-08.
- Konovalov, Alexander, and Zsolt Sándor. 2010. "On price equilibrium with multi-product firms." *Economic Theory*, 44(2): 271–292.
- Langer, Ashley. 2016. "(Dis)Incentives for Demographic Price Discrimination in the New Vehicle Market." Working Paper.
- Miller, Marty. 2019. "King of the Road: Breaking Down the Popularity of Pickup Trucks." https://www.experian.com/blogs/insights/2019/08/king-road-bre aking-popularity-pickup-trucks/, Last accessed on 2023-08-04.

- Murry, Charles. 2017. "Advertising in Vertical Relationships: An Equilibrium Model of the Automobile Industry." *Working Paper*.
- Murry, Charles, and Henry S. Schneider. 2016. "The economics of retail markets for new and used cars." In *Handbook on the Economics of Retail and Distribution*., ed. Emek Basker, 343–367. Edward Elgar.
- National Automobile Dealers Association. 2019a. "Driving Texas' Economy: Annual Contribution of Texas' New-Car Dealers." https://web.archive.org/we b/20190417192940/https://www.nada.org/WorkArea/DownloadAsset.aspx?i d=21474837361, Last accessed on 2023-06-11.
- National Automobile Dealers Association. 2019b. "NADA Data 2018: Annual Financial Profile of America's Franchised New-Car Dealerships." https://web.ar chive.org/web/20190417205114/https://www.nada.org/WorkArea/Download Asset.aspx?id=21474857318, Last accessed on 2023-06-11.
- Shiller, Benjamin Reed. 2020. "Approximating Purchase Propensities and Reservation Prices from Broad Consumer Tracking." International Economic Review, 61(2): 847–870.
- Shiller, Benjamin Reed. 2022. "Discreet Personalized Pricing." Brandeis University, Department of Economics and International Business School Working Papers 133.
- Thomadsen, Raphael. 2005. "The Effect of Ownership Structure on Prices in Geographically Differentiated Industries." *The RAND Journal of Economics*, 36(4): 908–929.
- Waldfogel, Joel. 2015. "First Degree Price Discrimination goes to School." The Journal of Industrial Economics, 63(4): 569–597.

- Wardlaw, Christian. 2019. "10 Most Popular Trucks." https://www.jdpower.co m/cars/trucks/10-most-popular-trucks, Last accessed on 2023-06-09.
- Wayland, Michael. 2021. "Pickup trucks dominate Americas 10 best-selling vehicles of 2020." https://www.cnbc.com/2021/01/08/pickup-trucks-dominate-ameri cas-10-best-selling-vehicles-of-2020.html, Last accessed on 2023-08-04.
- Ye, Junting, and Steven Skiena. 2019. "The Secret Lives of Names? Name Embeddings from Social Media."
- Ye, Junting, Shuchu Han, Yifan Hu, Baris Coskun, Meizhu Liu, Hong Qin, and Steven Skiena. 2017. "Nationality Classification Using Name Embeddings." *CIKM '17*, 18971906. New York, NY, USA:Association for Computing Machinery.
- Zettelmeyer, Florian, Fiona Scott Morton, and Jorge Silva-Risso. 2006. "Scarcity Rents in Car Retailing: Evidence from Inventory Fluctuations at Dealerships." National Bureau of Economic Research Working Paper 12177.

A Appendix for Results



Figure 9: Unconditional distribution of price sensitivities



Figure 10: Conditional distributions of price sensitivities