(Dis)Incentives for Demographic Price Discrimination in the New Vehicle Market*

Ashley Langer†

March 31, 2016

Abstract

New vehicles seem ripe for price discrimination based on consumer demographics with negotiated prices and substantial differences in preferences across groups, and yet the literature has found small and statistically insignificant differences in prices paid based on gender. This paper examines price discrimination in the new vehicle market based on gender and marital status and finds that while preferences do vary substantially across groups, full price discrimination based on gender and marital status would actually lower industry profits. This results from the fact that firms' ranking of demographic groups is asymmetric within vehicle segments in the way suggested by Corts (1998), where a firm's ability to increase its price to its "strong" group is undercut by its rivals' price decrease to that group. This asymmetry means that full price discrimination actually intensifies competition. I find that price discrimination based on gender and marital status would decrease industry variable profits by $6.4 billion per year, or over 3%. This suggests that laws that prohibit price discrimination based on certain demographics have the potential to decease price competition and increase firm profits by coordinating firms' commitment to uniform prices.

*This paper evolved from the first chapter of my dissertation. Amongst many others, I thank David Card, Gautam Gowrisankaran, Patrick Kline, Enrico Moretti, Aviv Nevo, Kenneth Train, Cliff Winston, Catherine Wolfram, Mo Xiao for their invaluable advice. Seminar participants at UC Berkeley, the San Francisco Federal Reserve, Collegio Carlo Alberto, the University of British Columbia, Arizona State University, Tufts University, the University of Toronto, Columbia University, the Ohio State University, the University of Wisconsin, Stanford University, the University of Arizona, the University of Chicago, the University of Maryland, and the University of Michigan provided valuable comments and discussion. Andrew Lumsden and Zachary Anderson were extraordinarily helpful with procuring the data.

†Department of Economics, University of Arizona. Email: alanger@email.arizona.edu

Introduction

Economists have long studied price discrimination based on individual's demographic characteristics in markets as varied as employment (Becker, 1957), cars (Ayres and Siegelman, 1995), housing (Yinger, 1998), fish (Graddy, 1995), trading cards (List, 2004), loans (Charles et al., 2008), and sex (Li et al., 2015). In particular, if consumers' preferences are correlated with their demographic characteristics and if prices are negotiated or otherwise set on a case-by-case basis, it seems plausible that firms could enact some form of price discrimination. This potential for price discrimination is important because it can lead to substantial reallocation of welfare between consumers and firms, but it is also of policy interest to those who worry about demographic groups being treated differently in markets based on protected characteristics like race and gender.

In particular, it would seem that the new car market is ripe for price discrimination based on consumers' demographics such as gender. Prices for new cars are negotiated between each consumer and salesmen who have a widespread reputation for "sizing people up" as they walk in the door. Additionally, preferences for new vehicles differ enough with gender that the terms "guy car" and "chick car" have entered the popular lexicon.[1] And yet a substantial literature has found economically small or statistically insignificant price differences for new cars across genders (e.g. Ayres and Siegelman (1995), Goldberg (1996), Harless and Hoffer (2002), and Scott Morton et al. (2003)). In fact, Goldberg (1996) explicitly states that the "striking feature" of her results is the lack of statistically significant price differences based on any socioeconomic characteristics.

There are at least three reasons why we might not observe price discrimination based on gender in the new vehicle market. First, it could be that preferences do not actually substantially differ by gender, and so there is just one optimal price for firms to charge both genders. Second, it could be that regulations, the threat of legal action, or other social constraints are substantial enough that firms would like to price discriminate absent legal ramifications but choose not to. Finally, it could be that firms' choose not to price discriminate based on gender because price discrimination would increase market competition and decrease firm profits as suggested by Corts (1998). This final outcome could occur if firms rank genders asymmetrically in the prices they would like to charge under full price discrimination and therefore each firm's price increase is undercut by its rivals' price decrease to the same group.

[1] e.g. http://www.cartalk.com/content/ultimate-guy-and-chick-cars-all-time

This paper explores this lack of price discrimination by gender observed in the new vehicle market. I begin by estimating the price elasticities of married and unmarried men and women in order to understand whether the previous literature's focus on average price differences across genders merely obscured price discrimination across vehicles that doesn't shift the average price paid overall. I then test whether the markups implied by full price discrimination fit the observed price variation better than alternative pricing models such as price discrimination by only some firms or price discrimination based on only marital status or gender but not both. Finally, I calculate optimal prices and variable profits (profits before the industry's substantial fixed costs are removed) in order to understand firms' incentives to price discriminate.

I estimate a random coefficients discrete choice demand model in the style of Berry et al. (1995), Berry et al. (2004), and Train and Winston (2007) for each demographic group separately and without imposing a supply-side moment. My estimates suggest that there are substantial and intuitive differences in demographic groups' demand for new vehicles. For instance, single people are generally more price sensitive than married people, married people are less interested in family un-friendly vehicles like pickup trucks, and single women are less interested in large, heavy vehicles than other demographic groups. The markups implied by these preferences similarly vary substantially over demographic groups and suggest that price discrimination would lead to substantially different prices to different demographic groups on identical vehicles.

I then use the preference estimates to calculate the markups firms would charge under a range of different approaches to price discrimination including uniform pricing, full price discrimination, price discrimination by only some firms, and price discrimination based on only one of the two demographic characteristics. I show that the variation in observed prices is better explained by firms not fully price discriminating, in fact I can reject the price discrimination model relative to every other pricing strategy. On the other hand, uniform pricing, price discriminating based only on marital status, and uniform pricing except for small subsets of firms are not rejected by any other model including full price discrimination. These tests suggest that even though preferences do vary substantially by gender, firms are not exploiting these differences in their observed prices.

While these tests show that firms' pricing strategies do not fit full price discrimination well, they do not shed much light on why this is true. In particular, regulations or the fear of legal ramifications could make uniform pricing more attractive to car-makers even if price discrimination would be profit-maximizing absent these constraints. However, Corts

(1998) also proposes an alternate explanation for why firms may not price discriminate even when they have the ability to. In certain situations, Corts argues, price discrimination may actually *increase* competition between firms and *decrease* profits. The argument is that if firms face an asymmetry in which demographic groups are considered "strong", thenone firm's price decrease to its weak group may undercut the price increase its rival would like to make to that same group enough to lower total industry profits.

Corts' model applies to two firms selling to two groups. The new car market is obviously substantially more complex with many firms selling to a wide variety of consumers. In order to better understand whether a Corts-style increase in competition could result from price discrimination in the new car market, I solve for the profit-maximizing prices that firms would charge under both uniform pricing and full price discrimination using the marginal costs implied by the uniform pricing model. I then compare the variable profit that firms would earn under each pricing strategy. I find that under full price discrimination based on gender and marital status, automotive industry profits are lower than under uniform pricing. Additionally, although there is no formal test of the type of asymmetry discussed by Corts in this complicated a market, I show that in many market segments (e.g. mid-size pickups) different firms would increase (and decrease) prices to different consumer groups under price discrimination. This suggests that increased competition coming from asymmetric "strong" and "weak" markets could be the potential mechanism through which price discrimination increases competition between firms and decreases industry profits.

If price discrimination based on gender and marital status would actually increase competition between firms and lower overall prices to consumers, then it raises questions about the role of government in enforcing anti-discrimination laws in the new vehicle market. Largely out of concerns about "taste-based" discrimination in the style of Becker (1957), laws exist to ensure that "protected classes" are not discriminated against in markets. While this paper in no way rules out that any taste-based discrimination is occurring or that price discrimination based on other demographics such as race would increase the prices paid by protected groups, these results do suggest that blanket policies outlawing demographic price discrimination may actually work in favor of firms by allowing them to collectively commit to avoiding profit-reducing price discrimination.

This paper contributes to the literature on price discrimination in automobiles by being the first to show that firms may prefer not to price discriminate even when consumer preferences differ based on observable demographics. This stands in contrast to the recent work by D'Haultfoeuille et al. (2014) which finds that firms in France do price discriminate based on income and age in the French automobile market, and that not price discriminating based on these demographics would decrease firm profits. Earlier work Verboven (1996) also showed that firms price discriminate across national borders in Europe. The results in this paper therefore show that there are at least two demographic characteristics upon which price discrimination would reduce firm profits, but there may be numerous dimensions along which price discrimination is still profitable. In particular, there is some evidence in Ayres and Siegelman (1995) that firms may be price discriminating by race, and this result is supported by Charles et al. (2008), who show that there seems to be price discrimination in loans by race. Therefore the generalizability of my results to other countries and demographic groups remains a question for further research.

The remainder of the paper is organized as follows. Section 1 explains the cross-sectional survey of new vehicle purchasers used in this analysis and compares that data to the previous literature's estimates of the difference in average prices paid by gender. Section 2 explains the empirical model used to estimate preferences and the approach to understanding which pricing strategy best fits the observed vehicle prices. Section 3 presents results of both the demand estimation and the comparison of various firm pricing strategies. It then also shows that price discrimination based on gender and marital status would reduce profits in this market and provides some evidence that the price changes under price discrimination follow the intuition presented in Corts (1998). Section 4 concludes.

1 Data

In order to understand firms' incentives to price discriminate, I use data on consumers' demographics, vehicles chosen (including the choice to not purchase a vehicle), prices paid, and vehicle attributes. I will discuss each of these before turning to my empirical model.

For information on consumers' demographics, vehicle choices, and prices paid, I pair a survey of new vehicle buyers from a major market research firm with data from the Current Population Survey. The survey of new vehicle buyers includes 25,875 respondents who purchased new vehicles in the second quarter of 2005, and includes self-reported information on the consumer's demographics, vehicle purchased, price paid, other vehicles considered, and whether the survey respondent is both the "principle buyer and driver" of the vehicle purchased. I limit my analysis to the 12,014 consumers who provide full demographic information, a price paid, and report being both the principle buyer and driver in order to assure that the demographic information observed by the dealer matches the preferences of the vehicle driver.

My analysis will focus on four demographic groups: married and unmarried men and women. These groups are large enough to estimate demand functions for each: 60.3% of my final sample is male and 68.4% is married. Gender and marital status are attractive groups to use for this analysis because they are fairly evenly distributed geographically, so it is likely that all dealerships have experience with consumers of all demographic groups. Additionally, gender is a readily observable variable to dealers and is often thought of as a dimension along which vehicle preferences may vary. Marital status may be less observable to dealers, so any differences in the amount of price discrimination based on marital status relative to gender might be related to consumers' ability to obscure their demographic group. Additionally, to the extent that married consumers are more likely than single consumers to be older and have larger households that potentially include children, I would expect married consumers' preferences to differ from single consumers of the same gender.[2]

In order to account for customers who decided not to purchase a new vehicle in the second quarter of 2005, I augment my data with information from the Current Population Survey on the total population of US adults in each demographic group. The objective is to construct a dataset that is representative of the population of potential new vehicle buyers and captures in the demographic heterogeneity in new vehicle sales of the survey data. To this end, I construct observations of "outside good" purchasers of each demographic group and add weights so that the full dataset reflects the population of potential new vehicle buyers.

The weights make use of two additional sources of information. First, I use data from the Automotive News Market Data Book on the total number of each vehicle model sold in the second quarter of 2005. I weight the survey responses in my data so that the total number of each model sold match the Automotive News number while the demographics reflect the survey responses. Additionally, I use information from GfK Automotive Research which says that approximately 20% of Americans considered buying a new vehicle in the previous year and assume that 10% of Americans considered buying a new vehicle in the second quarter of 2005. I weight the non-purchasing observations so that the total number of potential purchasers is equal to 10% of Americans and the overall demographic distribution

[2] Price discrimination based on race would be of particular interest to policy makers. Unfortunately, there are too few African American new vehicle purchasers in my sample to accurately estimate preferences. Additionally, I will need to assume that the outside good is identical for all demographic groups, which may not be true if African Americans are discriminated against in the used vehicle market as suggested by Charles, Hurst, and Stephens (2008).

matches that of the US population.[3]

In addition to this consumer choice data I use vehicle data from AutoData Solutions on the attributes of model year 2005 vehicles. This data includes extensive information on the vehicle, including the manufacturer's suggested retail price (MSRP), horsepower, curb weight, wheel base, fuel economy, turning radius, and whether the vehicle has stability control, traction control, or side airbags. This data is at the vehicle trim level, which allows it to differ for the same vehicle model based on differences such as engine type (e.g. V6 vs V8) or body style (e.g. hatchback vs sedan). Since my consumer choice data only specifies a consumer's purchase decision at the model level, I use the vehicle attributes of the trim with the lowest MSRP as the model attributes and consider any deviations from this to be captured by unobserved vehicle quality, which is allowed to differ across demographic groups. This reinforces the idea that consumers of different demographic groups might have different valuations of unobserved quality, since not only the vehicle's styling may be valued differently but also the average trim level chosen may vary by demographic group. To the extent that many options such as leather seats, rear spoilers, or sunroofs may be fairly inexpensive to produce but command a large markup, these options packages may be a way for firms to encourage consumers to self-select into options packages that are priced in order to further price discriminate.

Table 1 summarizes how the attributes of vehicles purchased in my sample vary across demographic groups. Married people generally purchase more expensive vehicles than single people and men purchase more expensive vehicles than women. Men purchase substantially more pickup trucks than women, while women purchase more SUVs than men (conditional on marital status). Single people are much more likely to purchase a car and much less likely to purchase a van relative to married people. Married men purchase vehicles with more horsepower than other groups, while single women purchase vehicles with less horsepower. Men generally purchase less fuel-efficient vehicles than women conditional on marital status, which may at least partially come from the fact that they purchase somewhat heavier vehicles than women. Finally, conditional on marital status men purchase vehicles with higher turning radius than women. These descriptive statistics suggest that there may be substantial differences in the preferences of these different demographic groups.

Finally, while my data includes the actual transaction price for the vehicle purchased, I clearly cannot observe the price that each consumer would have negotiated for every vehicle

[3] As a sensitivity check, I also construct a second set of weights that weight non-purchasers such that the full dataset has the demographic distribution of the survey of purchasers. This does not change results meaningfully.

Table 1: Descriptive Statistics of Demographic Groups' Vehicle Purchases

	Demographic Group			
Variable	Married	Married	Single	Single
Variable	Women	Men	Women	Men
Price Paid	30,732	32,758	25,569	29,625
Share Cars	0.51	0.48	0.69	0.61
Share SUVs	0.35	0.26	0.26	0.20
Share Pickups	0.04	0.17	0.02	0.15
Share Vans	0.09	0.08	0.03	0.04
Horsepower	199	217	174	198
Fuel Use (gal/100 mi)	4.49	4.86	4.04	4.50
Curb Weight (lbs)	3702	3882	3276	3565
Turning Radius (feet)	18.76	19.29	18.22	18.74
Number of Observations	2,852	5,366	1,920	1,876

Observations are weighted by the number of vehicles sold in the second quarter of 2005.

that she did not purchase. This is a long-standing problem in the vehicle-choice literature, to the extent that the nearly every discrete-choice vehicle demand paper uses the manufacturers' suggested retail price (e.g. Berry et al. (1995)) or an average transaction price (e.g. Allcott and Wozny (2013)) for all vehicles as the price the consumer actually pays.[4] When studying the variation in prices over consumers, this is clearly not a reasonable approach. Instead, I will use the average price paid by a consumer's demographic group for each vehicle in the choice set as the price the consumer pays for that vehicle. This assumes that the consumer is aware of the price paid for each vehicle by other consumers in her demographic group, and that she choses which vehicle to purchase based on those commonly known prices. Thus any actual variation in the price that a specific consumer would pay relative to other consumers in her demographic group must be random. While this is clearly a strong assumption, it avoids issues of heterogeneity in negotiation across individuals within a demographic group that are beyond the scope of this paper.

[4]The exception to this is the recent paper by (D'Haultfoeuille et al., 2014) who provide a method to recover price discriminating transaction prices from demographic-group market shares. However, they still assume that each demographic group pays a single price for each vehicle.

Comparison to the Literature

As noted, the previous literature on price discrimination in new vehicle sales finds economically small and frequently statistically insignificant differences in the average price of new vehicles by gender. This lack of substantial price differences by demographics appears in my data as well. Table 2 shows the simple differences in self-reported prices paid once vehicle fixed effects are controlled for. While men appear to pay slightly more for new vehicles and married customers pay slightly less, none of the demographic groups differ significantly in prices paid, and the point estimates of the differences are all less than $300 (conditional on marital status, the differences between genders are less than $200) . Given the average vehicle price of $30,639, these price differences are generally substantially below one percent of the vehicle transaction price. These results are very much in line with the literature discussed in the introduction (e.g. Ayres and Siegelman (1995), Goldberg (1996), Harless and Hoffer (2002), and Scott Morton et al. (2003)) who find small and generally statistically insignificant differences in purchases prices between genders. Furthermore, the model fixed effects explain approximately 74% of the variation in prices, but the demographics barely increase the R^2 of the regression at all. While this is far from conclusive evidence that firms are not price discriminating based on gender, it does suggest that this price discrimination is not substantial on average across vehicles. The full empirical model is therefore necessary to make more concrete statements about firms' incentives to price discriminate based on gender and marital status.

2 Empirical Model

In order to understand firms' incentives to price discriminate based on consumer demographics, the empirical approach proceeds in three steps. First, I estimate a random-coefficients discrete choice model of each demographic group's vehicle demand without imposing a supply-side model. Then I use these demand estimates to calculate firms' optimal markups under various pricing strategies. Finally, I calculate how profits would change if firms fully price discriminated.

2.1 Demographic Group Demand

I estimate demand completely separately for each demographic group using an estimation strategy that follows directly from Berry et al. (2004) but is estimated using maximum

Table 2: Correlations Between Transaction Prices and Demographics

	Dependent Variable Price Paid	
Male	8.36	-170.52
	(142.28)	(255.30)
Married	165.71	11.03
	(249.69)	(182.06)
Male and Married		270.13
		(291.50)
Vehicle Fixed Effects	Yes	Yes
N	12,014	12,014
Adjusted R^2	0.7391	0.7392

Standard errors are clustered by vehicle. Observations are weighted by the number of each vehicle sold in the second quarter of 2005. Results are qualitatively similar without weights, although the coefficient on male in column 1 does achieve statistical significance at the 5% level.

likelihood as in Train and Winston (2007) and omits any supply-side assumptions.

Consumers are each assumed to belong to a single demographic group, $d = 1, ..., D$. Within these demographic groups, consumers are heterogeneous along both observable and unobservable individual characteristics. Consumer i's utility for vehicle $j = 0, 1, ..., J$ is assumed to be:

$$U_{idj} = p_{jd}\tilde{\alpha}_{id} + \sum_k x_{jk}\tilde{\beta}_{idk} + \xi_{dj} + \varepsilon_{idj} \tag{1}$$

where p_{jd} is the price charged to i's demographic group d; $x_{j1}, ..., x_{jK}$ are the non-price attributes of vehicle j; ξ_{dj} is the preference of demographic group d for the unobservable attributes of vehicle j; and ε_{idj} is an extreme value type 1 residual preference parameter. The $\tilde{\alpha}_{id}$ and $\tilde{\beta}_{idk}$ are the individual's preference for vehicle attributes p_{jd} and x_k respectively, which are assumed to have the form:

$$\tilde{\alpha}_{id} = \bar{\alpha}_d + \sum_r z_{idr}\alpha^o_{dr} + \nu_{idp}\alpha^u_d \tag{2}$$

$$\tilde{\beta}_{idk} = \bar{\beta}_{dk} + \sum_r z_{idr}\beta^o_{dkr} + \nu_{idk}\beta^u_{dk}$$

Thus the individual's preference for vehicle attribute x_k is decomposed into a component ($\bar{\beta}_{dk}$) that is constant within that individual's demographic group, a component (β^o_{dkr}) that varies with the consumer's characteristics z_{idr} that are observed by the econometrician but not the dealer, and a component (β^u_{dk}) that varies with the consumer's characteristics ν_{idk} that are unobserved by the econometrician and the dealer, but are assumed to have a known distribution.

As noted, consumer demographics enter the model through both the d term that defines demographic groups and the z_{idr} term. The difference between these two components of preference is whether the demographic characteristic is observed by the dealer and therefore a potential basis for price discrimination. Since the objective of the model is to understand price discrimination based on demographics that are observable to the dealer, estimated preferences are allowed to be substantially more flexible in d than in z_{idr} (all coefficients may vary in d while only the preferences for price and certain vehicle characteristics may vary in z_{idr}). In particular, the unobservable quality of the vehicle, ξ_{dj}, varies in d but not in z_{idr}. Gender and marital status will be considered observable to the dealer, and thus enter as a part of d, while income will be considered unobservable to the dealer and enter as a part of z_{idr}.[5]

Combining equations (1) and (2) leads to the consumer's choice model:

$$U_{idj} = \delta_{dj} + \sum_r p_{jd} z_{idr} \alpha^o_{dr} + \sum_{k,r} x_{jk} z_{idr} \beta^o_{dkr} + p_{jd} \nu_{idp} \alpha^u_d + \sum_k x_{jk} \nu_{idk} \beta^u_{dk} + \varepsilon_{idj} \tag{3}$$

$$\text{where} \qquad \delta_{dj} = p_{jd} \bar{\alpha}_d + \sum_k x_{jk} \bar{\beta}_{dk} + \xi_{dj} \quad \text{for each} \quad j = 1, ..., J \tag{4}$$

The consumer chooses the vehicle $j = 1, ..., J$ or the outside option ($j = 0$, not purchasing a new vehicle) that maximizes this utility function. The outside option of not purchasing a new vehicle is modeled as having utility equal to $U_{id0} = \beta^o_d(\frac{1}{\text{Income}_i}) + \beta^u_{d0}\nu_{id0} + \varepsilon_{id0}$, where ν_i is a draw from a standard normal distribution. As this notation makes clear, there is a component (δ_{dj}) to each individual's utility for each vehicle that is common across all members of his or her demographic group d. Additionally, the term $\sum_r p_{jd} z_{idr} \alpha^o_{dr} + \sum_{kr} x_{jk} z_{idr} \beta^o_{dkr}$ allows consumers with different observable characteristics to have different tastes for certain vehicle attributes, and thus specifies the extent to which vehicle substitution varies with

[5] Demographics that are not directly observable to the dealer, z_{idr}, enter into price discrimination in expectation. Therefore, if married men are generally for higher income households than single women, the expected household income will enter into price discrimination but the variation in household income across married men will not.

observable consumer demographics. Finally, there is a component of consumer preference $(p_{jd}\nu_{idp}\alpha^u_{dk} + \sum_k x_{jk}\nu_{idk}\beta^u_{dk})$ that is unobserved by the econometrician but helps to explain why individuals may substitute more strongly between certain vehicles. The β^u_{dk} and α^u_d coefficients measure the standard deviation in the unobserved preference within demographic group d for the vehicle attribute conditional on the consumer's observed attributes. For notational ease, I define the vector of distributional coefficients $\theta_d \equiv [\alpha^o_{dr},\ \alpha^u_{dr},\ \beta^o_{dkr},\ \beta^u_{dkr}]'$.

I estimate the θ_d and δ_d coefficients via maximum-likelihood. Given the extreme-value error term, the probability of a consumer i in demographic group d choosing vehicle j given the vehicle attributes and consumer demographics can be expressed as an integral over the distribution of ν_{id}. Because the θ_d coefficients determine how consumers substitute between vehicles as attributes change, information on consumers' first and second choice vehicles aids identification of θ_d. Thus, the joint probability that consumer i chooses vehicle $j = 1$ out of the full choice set, and $j = 2$ out of the choice set with $j = 1$ and the outside good removed is:[6]

$$Pr_{i1}(\theta_d, \delta_d) Pr_{i2}(\theta_d, \delta_d|1) = \int \frac{\exp(V_{id1}(\nu_{id}; \theta_d, \delta_d))}{\sum_{l=0}^{J} \exp(V_{idl}(\nu_{id}; \theta_d, \delta_d))} \left(\frac{\exp(V_{id2}(\nu_{id}; \theta_d, \delta_d))}{\sum_{l=2}^{J} \exp(V_{idl}(\nu_{id}; \theta_d, \delta_d))} \right) f(\nu) d\nu$$

where $V_{idj}(\nu_{id}; \theta_d, \delta_d)$ is the non-stochastic component of consumer i's utility for vehicle j from equation (3). The log-likelihood function is then calculated as the sum of the log of this expected probability over consumers i in demographic group d with the integral approximated using simulation.

The log-likelihood function is maximized over θ_d. For each value of θ_d, I choose δ_d to set the predicted market shares for each demographic group equal to the observed market shares for that group as in Berry (1994):

$$\begin{aligned} S_{dj} &= \int_{z_{idr}} \int_{\nu} Pr_{idj}(\theta_d, \delta(\theta_d)) f(\nu) f(z_{idr}) d\nu dz_{idr} \\ &= Pr_{dj}(\theta_d, \delta(\theta_d)) \end{aligned} \tag{5}$$

where $f(z_{idr})$ is the pdf of the consumer characteristics z_{idr} in the demographic group d. Therefore it should be understood that the δ_d vector is estimated conditional on θ_d and is thus formally $\delta(\theta_d)$. The maximum-likelihood procedure solves for the value of θ_d that

[6] I remove the outside good from the second-choice choice set because the second choice information is based on the vehicle the consumer said she considered but did not purchase. It is not clear whether she would have purchased the second choice vehicle if the first choice were not available (she may not have purchased any vehicle), but it is her preferred alternative out of the set of vehicles once her first choice is removed.

maximizes the likelihood function subject to a market-share constraint that is a function of both θ_d and $\delta(\theta_d)$.

Once I have estimated θ_d and calculated $\delta(\theta_d)$, I use this mean preference vector to extract information about the $\bar{\alpha}_d$ and $\bar{\beta}_{dk}$ coefficients rather than just the θ_d coefficients. Recall from equation 4 that unobservable vehicle quality, ξ_{dj} may include vehicle attributes that allow firms to charge more for the vehicle. Therefore, an OLS regression of the δ_{dj} vector on vehicle price and attributes will estimate that consumers are less price sensitive than they actually are. In order to correct for this bias, I run a weighted IV regression of δ_{dj} on the vehicle price and attributes, using weights equal to the number of consumers in each demographic group who choose each vehicle. I use the standard Bresnahan (1987)/Berry et al. (1995) instruments plus a distance instrument used in Train and Winston (2007) and suggested by Gandhi and Houde (2015)Gandhi and Houde (2015):

$$\sum_{l \in f_j, l \neq j} x_{lk}, \qquad \sum_{l \in f_j, l \neq j} (x_{jk} - x_{lk})^2, \qquad \sum_{l \notin f_j} x_{lk}, \qquad and \qquad \sum_{l \notin f_j} (x_{jk} - x_{lk})^2 \tag{6}$$

which are the sum of each vehicle attribute for competing vehicles produced by the same firm as vehicle j, f_j, the sum of each vehicle attribute for competing vehicles produced by other firms, and the sum of the squared distance in attribute space between the vehicle and all others sold by the same firm and all others sold by other firms. These instruments are intended to capture the extent of price competition faced by vehicle j in the market. For instance, if a vehicle is competing with a set of vehicles that have particularly high horsepower, then competitive pressure will keep the vehicle's price fairly low conditional on its attributes. If the observed price is actually high conditional on attributes, it must be that the vehicle has a high level of of unobservable quality that is increasing its demand. As with the rest of the demand estimation, the instrumental variable regression is run completely separately for each demographic group in order to allow the competitive pressure on price created by competing vehicles' attributes to vary over demographic groups.

2.2 Understanding Firms' Incentives to Price Discriminate

Once the demographic groups' preference parameters (θ) are estimated, I can use these parameters in combination with the assumption that firms are playing a Bertrand-Nash equilibrium to understand whether firms are price discriminating. In particular, I assume that firms maximize profits by solving the vector of well-known first-order conditions:

$$p_{dj} = c_j - \Sigma(p_{dj}|\theta)^{-1} Q(p_{dj}|\theta) \tag{7}$$

where $\Sigma(p_{dj}|\theta)^{-1}$ is the matrix of derivatives of purchase quantities with respect to prices for vehicles sold by the same manufacturer (and 0 otherwise), and $Q(p_{dj}|\theta)$ is the predicted sales of each vehicle. Both Σ and Q are functions of the observed prices and the estimated preference coefficients, θ. For ease of exposition, I will use $\hat{M}_{dj}(p_{dj}|\theta)$ to represent the markup on product j to demographic group d, or $\hat{M}_{dj}(p_{dj}|\theta) \equiv -\Sigma(p_{dj}|\theta)^{-1} Q(p_{dj}|\theta)$.

These markups can be calculated under different assumptions about the pricing strategy of the firms. For instance, firms could be fully price discriminating (charging a distinct price on each vehicle to each demographic group) or be pricing uniformly on each vehicle to all demographic groups. In practice, I calculate the markups that firms would charge under both of these strategies as well as all firms only price discriminating on gender or marital status but not both and firms in each region (US, Europe, Japan, and Korea) price discriminating while all others price uniformly.

Under the assumption that the marginal cost of a vehicle does not vary with the demographic group that purchases it, the variation in the observed prices across vehicles and demographic groups should be best explained by the true pricing strategy's calculated markups and vehicle fixed effects.Therefore, I preform a Vuong goodness-of-fit test (Vuong, 1989) of the model:

$$p_{dj} = \gamma_j + \hat{M}^s_{dj}(p_{dj}|\theta) + \varepsilon_{dj} \tag{8}$$

where γ_j are estimates of the marginal cost of each vehicle, the s superscript on the markup indexes each pricing strategy, and the ε_{dj} captures measurement error in the observed prices. The Vuong tests are bi-lateral tests between models, so while it may not be possible to identify the firm pricing strategy that is statistically significantly better than every other strategy at explaining the variation in observed prices, I may be able to reject that certain models are best at explaining observed prices and to observe tendencies for certain types of models to explain the data better than others.

Finally, once I have identified a pricing strategy that best explains the variation in the observed prices, I can calculate how variable profits would change under more or less price discrimination. To do this I use the estimates of marginal costs from the estimation of equation 8 that best fits the data and then find the vector of prices that solve the set of first order conditions in equation 7 under various pricing strategies. This calculation is different than the approach to comparing pricing strategies to the observed data because the regressions in equation 8 are only valid under the assumption that the markups are generated by firms' actual pricing strategy. In order to calculate counterfactual variable profits, I solve for the full counterfactual vector of prices conditional on the estimated marginal costs and the counterfactual pricing strategies.

3 Results

3.1 Demand Estimation Results

In order to understand firms' pricing incentives, it is first critical to understand how demographic groups vary in their preferences for new vehicles. I will present the demand estimation results for all of the demographic groups by first looking at the mean preference coefficients, then the estimates of preference heterogeneity within demographic groups, and finally the price elasticities and fully price-discriminating markups predicted by my estimates.

Regressing the δ_d vectors on vehicle attributes using weighted instrumental variables generates the mean preference coefficients for each demographic group.In these regressions, I include price (instrumented with the Berry et al. (1995) instruments as discussed earlier), whether the vehicle is a car, and if so whether it is a "sporty" car (generally a small, high horsepower car like the Acura RSX, the Mazda Miata, or the VW GTI), and whether the vehicle is a truck. I also include vehicle curb weight, number of passengers, turning radius, and whether the vehicle is manufactured by a domestic manufacturer in the mean preference specification.

Table 3 reports the mean preference coefficients for each demographic group. Married men are the least price sensitive on average while single women are the most price sensitive. All four groups dislike cars relative to the omitted vehicle types (SUVs and vans), while only married consumers dislike pickup trucks relative to the omitted category of SUVs and Vans. All groups except single men dislike sporty cars even more than they dislike cars as a whole (the total preference for sporty cars relative to SUVs and Vans is the sum of the car coefficient plus the sporty car coefficient). The preferences for vehicle characteristics are generally what we might expect. All demographic groups except for single women like bigger vehicles, as measured by curbweight, which are also generally safer. Conditional on vehicle weight and type, however, consumers dislike vehicles with higher seating capacity, a distaste which is stronger for men than women. All consumers prefer vehicles with smaller turning

Table 3: Mean Preference Coefficients by Gender and Marital Stauts

	Demographic Group			
Variable	Married Women	Married Men	Single Women	Single Men
Price	-1.40***	-1.30***	-1.81***	-1.69***
(tens of thousands of dollars)	(0.32)	(0.35)	(0.31)	(0.32)
Car	-1.80***	-1.80***	-1.66***	-0.89**
	(0.34)	(0.37)	(0.38)	(0.37)
Pickup	-3.06***	-0.86*	-0.37	0.24
	(0.41)	(0.46)	(0.44)	(0.44)
Sporty Car	-1.14***	-1.59***	-1.00**	-0.53
	(0.44)	(0.49)	(0.48)	(0.50)
Curbweight	0.94***	1.07***	-0.02	1.52**
(thousands of pounds)	(0.36)	(0.41)	(0.46)	(0.40)
Number of Passengers	-0.18*	-0.57***	-0.20	-0.53***
	(0.11)	(0.12)	(0.12)	(0.12)
Turning Radius	-4.52***	-4.34***	-3.19***	-5.08***
(feet)	(0.44)	(0.48)	(0.49)	(0.49)
Domestic	-0.69***	-0.46*	-0.89***	-0.53**
	(0.23)	(0.25)	(0.25)	(0.25)
Number of Observations	213	213	213	213

Instrumental variables regression of the mean preference of each group for each vehicle on vehicle characteristics. Instruments are functions of the vehicle attributes of competing vehicles, as discussed in the empirical strategy section. Weighted instrumental variables standard errors in parentheses, where the weights are equal to the number of observations for that demographic group that purchased that vehicle in the maximum likelihood stage. Significance level indicated by: *=10%, **=5%, ***=1%.

radii, which likely captures other aspects of the vehicle's performance as well. This distaste for poor performance is particularly strong for single men and much weaker for single women. Finally, all demographic groups prefer imported vehicles to domestic vehicles controlling for attributes, but this preference is particularly strong for women.

The specification of consumer demand heterogeneity includes price, price divided by household income, the four vehicle types (car, truck, SUV, and van) as well as the fuel use (in gallons per mile), horsepower, and curbweight, an indicator variable for the outside good and the outside good divided by household income. I specify all vehicle characteristics that are not interacted with observable consumer characteristics as having normally distributed unobservable heterogeneity. The coefficients on all of the normally distributed unobservable heterogeneity terms can be interpreted as the standard deviation in the demographic group's preference for the vehicle attribute, while the coefficients on price divided by income and the outside good divided by income are the extent to which the preference for price and the outside good vary with income.

Table 4 presents the coefficient estimates for these consumer heterogeneity terms by demographic group. The first thing to notice is that there is very little heterogeneity in the price coefficient within a demographic group once household income is taken into account, but household income dramatically affects consumers' sensitivity to a vehicle's price. The sign of this effect is what would be expected: consumers in lower income households react more negatively to a vehicle's price than consumers in higher income households. Women's price sensitivity appears to be somewhat more affected by household income than men's. There is substantial variation in preferences within groups for vehicle types (SUV, truck, van, car) for all demographic groups except single women, although single women do show variation in the strength of their preference for cars, and the variation in their preference for vans is large but not statistically significant. All groups show some evidence of heterogeneity in their preference for horsepower, although this heterogeneity is larger for women and not statistically significant for single men. The heterogeneity in preference for fuel use is only statistically significant for married men, while the heterogeneity in preference for curb weight is only statistically significant for single women. Recall that single women were the only group to not have a strong preference for vehicle weight on average, so it makes sense that this is the result of some single women preferring heavier vehicles and some disliking heavier vehicles.

Finally, demographic groups generally do not have large variation based on unobservable attributes in their preference for the outside good, but their preference for the outside good

Table 4: Preference Heterogeneity Coefficients by Gender and Marital Status

Variable	Variable Type	Married Women	Married Men	Single Women	Single Men
		Demographic Group			
Price	Std Dev	0.08**	0.05	0.01	0.002
(tens of thousands of dollars)		(0.04)	(0.03)	(0.06)	(0.048)
	Divided by Income	-7.78***	-5.68***	-6.75***	-5.25***
		(1.33)	(0.84)	(1.52)	(1.29)
SUV	Std Dev	1.46***	2.29***	0.04	1.98***
		(0.53)	(0.38)	(0.24)	(0.72)
Pickup	Std Dev	2.73***	2.98***	0.24	2.22***
		(0.72)	(0.45)	(0.42)	(0.84)
Van	Std Dev	1.72***	1.90***	2.17	1.97
		(0.67)	(0.54)	(1.41)	(1.41)
Car	Std Dev	2.80***	3.40***	2.83***	2.87***
		(0.48)	(0.45)	(0.66)	(0.78)
Horsepower	Std Dev	0.72***	0.36*	0.70*	0.50
(hundreds)		(0.23)	(0.20)	(0.41)	(0.31)
Fuel Use	Std Dev	0.13	0.27***	0.16	0.07
(gal./hundred miles)		(0.10)	(0.06)	(0.14)	(0.18)
Curb Weight	Std Dev	0.14	0.09	0.58 *	0.02
(thousands of pounds)		(0.13)	(0.14)	(0.32)	(0.03)
Outside Good	Std Dev	0.72	0.20	0.21	0.21
		(0.67)	(0.61)	(0.79)	(0.79)
	Divided by Income	-17.44***	-11.23***	-14.14***	-12.62***
		(3.51)	(2.43)	(3.87)	(3.62)
Number of Halton Draws		200	200	200	200
Number of Observations		3092	5606	2400	2356

Standard errors in parentheses. Significance level indicated by: *=10%, **=5%, ***=1%.
Coefficients estimated with maximum likelihood.

does vary with income. Perhaps the most surprising coefficient in the demand estimation results is that consumers from households with lower income have a stronger distaste for the outside good (which, recall, includes purchasing a used vehicle or continuing to drive your current vehicle). This result is less surprising in light of the fact that the price of the outside good has already been controlled for in the price interacted with household income coefficient, so this result must be picking up differences in preferences for the attributes of new cars relative to the outside good or differences in the methods of financing new vehicle purchases. If low income households are using leases to purchase new vehicles while high income households purchase new vehicles with cash, then we might expect a stronger preference for the outside good among high income households as I find here.

The estimated preferences generate reasonable distributions of implied price elasticities and price discriminating markups. Table 5 shows that elasticities range from approximately -2 to -6 with mean elasticities in the -3.5 to -4.5 range for all demographic groups. These are very much in line with results in other contexts (Berry et al. (1995), Whitefoot et al. (2013)). Predicted markups if each demographic group faces their own unique (price discriminating) price are similarly reasonable and range from approximately \$4,500 to nearly \$14,000. The average predicted markups range from 24% of purchase price (for single women) to 32% of purchase price (for single men), which is in line with the 31.5% used by the US Environmental Protection Agency to estimate the impact of emissions control regulations (Rogozhin et al., 2009). We see relatively high markups relative to transaction prices in the new vehicle market because the high fixed costs for each model produced prevent firms from bringing vehicles to the market that aren't expected to earn high markups over marginal costs.

Table 5 suggests that there are substantial differences in the price discriminating markups across groups. Just looking at the average markups, the optimal price discriminating markup for married men is \$2,258 or 33.7% higher than the average markup for single men. In order to compare variation in the predicted markups to the variation in the average prices by demographic group, I regress each on vehicle fixed effects and dummies for gender and marital status, weighting by the number of consumers of each demographic group who purchase each vehicle. I then compare these coefficients to a regression of the average purchase price for each vehicle to each demographic group on vehicle fixed effects and demographic dummies, similarly weighted. Table 6 presents these results. The first thing to notice is that the differences in markups by gender and marital status are substantially larger than the differences in prices. The results suggest that men pay \$65 more than women on average and married people pay \$76 less than single people, but neither of these results is statistically

Table 5: Elasticities and Price Discriminating Markups

		Demographic Group			
		Married Women	Married Men	Single Women	Single Men
Elasticity					
	Min	-5.81	-6.26	-5.07	-7.66
	Max	-2.46	-1.90	-2.77	-2.28
	Mean	-3.66	-3.51	-4.13	-4.56
Implied Markup ($)					
	Min	4,995	6,466	4,510	4,885
	Max	12,487	12,869	13,758	9,329
	Mean	8,443	8,950	7,520	6,692

The mean elasticity and markup are the unweighted mean over the vehicles in the choice set.

significant.[7] However, men's average predicted markup is $245 higher than women's and married consumers' markups are $1,877 higher than single consumers'. Further, gender and marital status do not help to explain the variation in average prices at all: the adjusted R^2 of a regression of prices on just vehicle fixed effects is 0.9809 while the adjusted R^2 of the regression that adds gender and marital status is 0.9808. However, these demographics do substantially help to explain the variation in the markups: the adjusted R^2 of a regression on just vehicle fixed effects is 0.5950 while the adjusted R^2 of the regression that adds gender and marital status is 0.9240. These regressions suggest that predicted markups are substantially more correlated with gender and marital status than the average prices, which at least suggests that firms may not be fully price discriminating.

3.2 (Dis)Incentives for Price Discrimination

With the demographic-group-specific demand estimates, I can think about firms' incentives to price discriminate in the market for new vehicles. As discussed above, the predicted markups for each demographic group under full price discrimination are very much in line with estimates used in other contexts. However, this does not mean that firms are engaging in full price discrimination or that it is optimal for them to do so. As explained in section

[7] The effect of demographics on prices presented in Table 6 differ somewhat from those presented in Table 2 because the unit of observation is the vehicle-demographic group rather than the individual purchase.

Table 6: Variation in Average Prices and Predicted Markups

	Dependent Variable	
	Average Price	Predicted Markup
Male	64.87	245.43
	(168.15)	(45.40)
Married	-75.95	1,877.44
	(127.78)	(67.86)
Vehicle FEs	Yes	Yes
Adjusted R^2	0.9808	0.9240
Adjusted R^2 w/o demographics	0.9809	0.5950

Standard errors are clustered by vehicle. Results are qualitatively similar if the standard errors are unclustered. All regressions are weighted using the number of vehicles purchased by each group.

2.2, I calculate the optimal markups for firms to charge under different pricing strategies and then conduct a Vuong test (Vuong, 1989) to understand which set of markups best explains the observed price variation.[8]

The Vuong tests are bi-lateral tests of model fit based on a comparison of log likelihood values. In practice this means that I cannot reject all models in favor of a single pricing strategy, but clear patterns emerge from the test statistics. Table 7 presents the test statistics for a selection of important pricing strategies. The statistics in the table are for the test of the column model relative to the row model, so a positive test statistic favors the column label while a negative test statistics favors the row label. Full price discrimination is rejected relative to every other pricing strategy tested, so any of the tested markups explain the variation in the observed prices better than the fully price discriminating markups. On the other hand, fully uniform pricing is not rejected relative to any of the other pricing strategies tested. This lack of rejection is also true for the model where prices are uniform except for GM price discriminating, the model of uniform prices except for Korean firms, and the model where all firms price discriminate on marital status but not gender.[9] These

[8]I test 15 different pricing strategies: uniform pricing, price discrimination by only US firms, Japanese firms, European firms, or Korean firms, price discrimination by only Chrysler, Ford, GM, Toyota, Honda, or Nissan, full price discrimination, or price discrimination based on only gender or marital status but not both.

[9]These same results hold although with somewhat less power for a model which allows for "taste-based" discrimination in the style of Becker (1957) where firms increase prices by fixed amounts on all vehicles for each demographic group, e.g. where $p_{dj} = \gamma_j + \gamma_d + \hat{M}^s_{dj}(p_{dj}|\theta) + \varepsilon_{dj}$.

tests overwhelmingly suggest that firms are not fully price discriminating and may, in fact, not be price discriminating based on gender at all.

Table 7: Vuong Test Statistics for Selected Pricing Strategies

Pricing Strategy	Uniform	Uniform Except US	Uniform Except Japan	Full Price Discrim.	Price Discrim. on Gender
Uniform					
Uniform Except US	1.41				
Uniform Except Japan	3.27	0.17			
Full Price Discrimination	3.30	3.30	2.18		
Price Discrimination on Gender	2.57	-0.33	-0.95	-2.70	
Price Discrimination on Marital Status	0.12	-1.36	-1.99	-3.68	-1.31

Entries are the z statistic for the test that the top pricing model performs better than the pricing model on the left hand side. Price discrimination on gender and price discrimination on marital status are the pricing strategies to fully price discriminate based on gender or marital status *only.* Test statistics for all pricing strategies tested are available upon request.

Given that more uniform approaches to pricing are better at explaining the variation in observed prices than full price discrimination and at least as good as any other model, I use the marginal costs implied by the uniform pricing model and calculate firms' variable profits under uniform pricing relative full price discrimination. The change in variable profits incurred by moving from uniform pricing to full price discrimination in the second quarter of 2005 are presented in Table 8. These results use the vehicle marginal costs implied by a current uniform pricing strategy but solve for the profit-maximizing prices for firms to charge under both the uniform and full price discrimination pricing strategy given demographic group preferences. Under full price discrimination, variable profits to the industry as a whole would have been 1.6 billion dollars lower than under uniform pricing. Most of this difference comes from the fact that the "Big 3" US manufacturers plus Honda and Toyota would have had substantially lower variable profits under full price discrimination. However, some smaller firms such as Daimler, Nissan, and VW would have had higher variable profits if the entire industry had price discriminated than under uniform pricing.

In Corts (1998), firms may want to avoid price discrimination if they are "asymmetric" in that firms offer products that appeal to different "strong" markets. While this asymmetry is much easier to observe in an industry where there are two firms selling to two markets, I do find evidence of asymmetry within market segments. For instance, in the mid-sized pickup truck market (e.g. Chevrolet Colorado or Ford Ranger), firms vary in how they order demographic groups with respect to optimal prices to charge under price discrimination. All firms would like to charge married women the highest price (likely a result of their low

Table 8: Change in Variable Profit with Full Price Discrimination

Firm	Millions of 2005 $
BMW	-3.2
Chrysler	-174.8
Daimler	2.1
Ford	-290.6
GM	-698.7
Honda	-306.7
Hyundai	0.6
Isuzu	-1.8
Kia	-11.0
Mitsubishi	3.5
Nissan	22.1
Porsche	2.5
Subaru	-11.1
Suzuki	1.4
Toyota	-142.1
VW	7.8
Total	-1,600

Total industry variable profits in Q2 of 2005 are $49.9 billion under uniform pricing and $48.3 billion under full price discrimination.

price sensitivity), but after that Ford charges higher prices to single women than married men while other manufacturers do the reverse. In fact, Nissan and Chrysler both would charge single women the lowest price under price discrimination, while GM and Toyota would charge single men the lowest price and Ford would charge married men the lowest price. These substantial differences in the price ordering of demographic groups under price discrimination is observable in many market segments where each firm produces relatively few vehicle models. In larger segments (e.g. midsize cars) there are differences in the ordering of demographic groups both across and within firms, so it is more difficult to find clear evidence of asymmetry. While the most common ordering is married women > married men > single men > single women, nearly 40% of vehicles would have a different ordering of prices under full price discrimination, suggesting substantial potential for asymmetry.

Finally, as suggested in Corts (1998), it may be difficult for firms to coordinate on uniform pricing if there are benefits to individual firms to price discriminating while their rivals price discriminate. My results suggest that firms are maintaining coordination of uniform pricing based on gender (and, to a lesser extent, marital status) in the new vehicle market. One could imagine that this ability to coordinate is the result of laws aimed at preventing price discrimination based on gender, although I have no way of testing that in my data. However, my results do suggest that the automotive industry as a whole should not be opposed to laws that outlaw price discrimination based on gender, since that reduces price competition between firms and increases industry profits.

4 Conclusion

This paper suggests that the lack of price discrimination based on gender noted in the earlier literature is likely the result of the fact that firms' product offerings mean that variable profits are lower under price discrimination than under uniform pricing. This result stands in contrast to recent work by D'Haultfoeuille et al. (2014) that finds that price discrimination based on age and income is profitable for vehicle manufacturers in France. The different results in this context could stem from the fact hat automotive firms in France do not face the same "asymmetry" in their pricing to different income and age groups that I observe in the U.S. based on gender and marital status. This explanation particularly makes sense for income, in that higher income groups' lower price sensitivity might mean that all firms prefer to charge high income groups higher prices. The differing results for different demographic groups makes it clear that my results should not be interpreted as evidence that there is no price discrimination occurring in the U.S. market, but rather that price discrimination based on gender and marital status appears to increase price competition and decrease firm profits. This result explains the lack of observed price differences by gender that have been found in the earlier literature in the U.S., and the methodology in this paper could be employed to understand firms' incentives to price discriminate based on other demographic characteristics.

One key discussion in Corts (1998) is the fact that it is often difficult for firms to commit to not price discriminating since firms can often increase profits if they are the only firm price discriminating while their rivals price uniformly. In the context of potential price discrimination based on consumer demographics, public policy may actually play a role in helping firms to commit to uniform pricing. Public discourse may argue that legal and social pressure is necessary to prevent firms from price discriminating and to "protect" certain groups from over-charging. However, this same pressure may allow firms to jointly commit to uniform pricing, which decreases price competition and increases firm profits. To my knowledge, this is the first paper to empirically show this somewhat counter-intuitive result that anti-discrimination laws could increase firm profits.

References

Allcott, Hunt and Nathan Wozny, "Gasoline Prices, Fuel Economy, and the Energy Paradox," *Review of Economics and Statistics*, October 2013, *96* (5), 779–795.

Ayres, Ian and Peter Siegelman, "Race and Gender Discrimination in Bargaining for a New Car," *The American Economic Review*, 1995, *85* (3), 304–321.

Becker, Gary S., *The Economics of Discrimination*, Chicago: University of Chicago Press, 1957.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63* (4), 841–890.

_ , _ , and _ , "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market," *Journal of Political Economy*, 2004, *112* (1), 68–105.

Berry, Steven T., "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 1994, *25* (2), 242–262.

Bresnahan, Timothy F., "Competition and Collusion in the American Automobile Industry: The 1955 Price War," *The Journal of Industrial Economics*, 1987, *35* (4), 457–482.

Charles, Kerwin Kofi, Erik Hurst, and Melvin Stephens, "Rates for Vehicle Loans: Race and Loan Source," *The American Economic Review*, 2008, *98* (2), 315–320.

Corts, Kenneth S., "Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment," *The RAND Journal of Economics*, 1998, *29* (2), 306–323.

D'Haultfoeuille, Xavier, Isis Durrmeyer, and Philippe Fevrier, "Automobile Prices in Market Equilibrium with Unobserved Price Discrimination," 2014.

Gandhi, Amit and Jean-Francois Houde, "Measuring Substitution Patterns in Differentiated Products Industries: The Missing Instruments," 2015.

Goldberg, Pinelopi Koujianou, "Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 1996, *104* (3), 622–654.

Graddy, Kathryn, "Testing for Imperfect Competition at the Fulton Fish Market," *The RAND Journal of Economics*, 1995, *26* (1), 75–92.

Harless, David W. and George E. Hoffer, "Do Women Pay More for New Vehicles? Evidence from Transaction Price Data," *American Economic Review*, 2002, *92* (1), 270–279.

Li, Huailu, Kevin Lang, and Kaiwen Leong, "Does Competition Eliminate Discrimination? Evidence from the Commercial Sex Market in Singapore," Working Paper 20911, National Bureau of Economic Research January 2015.

List, John A., "The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field," *The Quarterly Journal of Economics*, 2004, *119* (1), 49–89.

Morton, Fiona Scott, Florian Zettelmeyer, and Jorge Silva-Risso, "Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities?," *Quantitative Marketing and Economics*, March 2003, *1* (1), 65–92.

Rogozhin, Alex, Michael Gallaher, and Walter McManus, "Automobile Industry Retail Price Equivalent and Indirect Cost Multipliers," Technical Report EPA-420-R-09-003 2009.

Train, Kenneth E. and Clifford Winston, "Vehicle Choice Behavior and the Declining Market Share of U.s. Automakers*," *International Economic Review*, November 2007, *48* (4), 1469–1496.

Verboven, Frank, "International Price Discrimination in the European Car Market," *The RAND Journal of Economics*, 1996, *27* (2), 240–268.

Vuong, Quang H., "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses," *Econometrica*, 1989, *57* (2), 307–333.

Whitefoot, Kate S., Meredith Fowlie, and Steven J. Skerlos, "Compliance by Design: Industry Response to Energy Efficiency Standards," 2013.

Yinger, John, "Evidence on Discrimination in Consumer Markets," *The Journal of Economic Perspectives*, 1998, *12* (2), 23–40.