The Gender Gap in Price Negotiations: Evidence from New Car Sales*

Ambarish Chandra Sumeet Gulati James Sallee

December 6, 2013

Abstract

Increases in the educational attainment and labor force participation of women have constituted some of the biggest societal changes in North America over the past few decades. Yet, economic research into the consequences of this phenomenon has been restricted almost exclusively to the labor market. We investigate whether women's performance in price negotiations has improved relative to that of men. Using detailed transaction level data from the automobile industry, we show that the age premium rises more steeply for women than for men, consistent with the hypothesis that higher levels of education and work experience have led women to close the gender gap in this industry.

Keywords: Gender; Age; Automobiles; Negotiations; Bargaining. **JEL Codes:** J14, J16, L62.

*Chandra: Department of Economics, University at Albany, ambarishchandra@gmail.com; Gulati: Food and Resource Economics, University of British Columbia, sgulati@mail.ubc.ca; Sallee: Harris School of Public Policy, University of Chicago, sallee@uchicago.edu

1 Introduction

The last several decades have seen dramatic changes in the educational attainment and labor force participation of women. In North America, women now make up about 45% of the labor force, and almost 60% of college graduates, a stark change from the 1970s. This phenomenon has had huge societal effects, such as on the rate and age of marriage, as well as on fertility rates. It has had equally profound economic impacts, such as on wages and career opportunities for women.

Most prior economic research into the consequences of increased women's education and employment has focused on the labor market itself. Economists have documented the effects on women's wages, most notably the fact that the gender wage gap has narrowed but not closed.[1] Various studies have also examined changes in female representation in managerial positions.[2] There has, however, been remarkably little research on the impact of these societal changes in other markets. In particular, few prior studies have examined whether there have been changes in the prices that women pay in important transactions that involve negotiations, other than studies of wage negotiations.[3]

Personal interaction and negotiation influence the prices or quality of products in several key markets. We negotiate with service providers, financial institutions and, most importantly, in the markets for housing and automobiles. Individual characteristics can potentially play an important role in negotiated outcomes. In this paper, we examine whether prices negotiated by women have shown significant improvements over time. In particular, we ask whether older women today, who on average have lower education levels and had a lower likelihood of having been employed when in their 20s and 30s, fare worse in price negotiations than younger women. We examine how each cohort of women performs relative to men of their same age, treating these men as a control group since they did not undergo the same transformative experiences that women did.

It is challenging to identify appropriate markets for this study. The ideal market would provide data on both negotiated prices and consumer demographics, and would also involve transactions of identical goods, or else enough data on varying characteristics to control for systematic differences in the choice of goods by different types of consumers. However, prices are generally not open to negotiation in many everyday markets where the personal characteristics of consumers are observed. Conversely, it is difficult to obtain price data in markets with negotiated prices since these usually occur in private or informal settings; for example, when bargaining for loans with financial institutions, selling goods such as used cars via Craigslist, or hiring contractors to perform residential services. The housing market appears promising, but it is rare to find housing transactions that involve identical units — a sample of housing transactions will always involve considerable heterogeneity in characteristics that are difficult to control for precisely.

Therefore, we look to the new car market, which we believe offers the ideal setting for our study for a number of reasons. First, final prices in this market are commonly negotiated

[1] See, for example, Blau and Kahn (2007) and O'Neill (2003).

[2] See Bertrand and Hallock (2001) and Bertrand et al. (2010).

[3] Specifically, we are not aware of prior studies of this issue in the Economics literature. Other fields, particularly Psychology, have investigated this topic, though primarily in laboratory settings. See Stuhlmacher and Walters (1999) for a survey.

and there can be differences of hundreds of dollars in the amounts that different consumers pay for the same new car. Second, this market involves the sale of many identical goods: knowing the make, model and trim of a vehicle pins down almost exactly the type of good that different consumers buy. Other differences in transactions, such as the location of car dealers and the timing of the transaction, can easily be controlled for. Finally, we have data on two key elements of new car transactions: the final price paid by consumers, and the dealer's invoice price, which represents the opportunity cost to the dealer of selling the car, and therefore includes the cost of unobserved options purchased by consumers. Therefore we can construct precise measures of the dealer's profit margin — which is rare in almost any industry — and examine whether these margins are systematically different for different types of consumers.

Previous research into the automobile industry has, of course, examined the role of gender and other demographics in price negotiations. Famously, Ayres and Siegelman (1995) used an audit study and found that women and minorities are disadvantaged in the new car market, as both groups are offered higher initial prices by dealers as well as negotiate higher final prices. Since then, a number of studies have re-examined the issue of gender differences in the new car market, including Goldberg (1996) and Harless and Hoffer (2002). Both those studies cast doubt on the finding that women pay more for new cars. More recently, Morton et al. (2003) show that, while minority customers pay a higher price than others, this can be explained by their lower access to search and referral services. Similarly, Busse et al. (2013) find that women are quoted higher vehicle repair prices than men when callers signal that they are uninformed about prices, but these differences disappear when callers mention an expected price for the repair.

Importantly, no previous study has examined the interaction of age and gender in the new car market, especially in the context of possibly different behavior exhibited by different cohorts.[4] Our study does so explicitly, and also makes other contributions to the existing literature. First, we use a dataset involving more than half a million transactions in the new car market, a far larger sample than was available in the earlier studies described above. Second, we have data on both final transaction prices and dealers' invoice prices, which allows us to calculate and study dealer margins. Third, the large sample allows us to divide consumers into fine age and gender categories, and also to use a large set of fixed-effects that control precisely for vehicle characteristics.[5]

Following Busse et al. (2013), we estimate equations for transaction prices and dealer margins controlling for a range of demand and supply covariates. Our data come from JD Power, and are drawn from a very large set of new car transactions in Canada between 2004 and 2009. Benefiting from the large size of our sample, we identify variations in dealer margins within model, year, province and trim variations, and across several gender and age categories. Our estimating strategy minimizes the potential impact of other influences: whether the vehicle was leased or financed; whether it was purchased at the end of the month

[4] Harless and Hoffer (2002) do examine customer age, finding that older consumers generally pay more on average, but do not examine the relative performance of each gender across age cohorts. Langer (2011) examines the interaction between the gender and marital status of consumers, but does not focus on cohort effects either.

[5] This also allows us to use simple methods with confidence, as we know both the trim level and the cost of the vehicle, in contrast to, say, Langer (2011).

or year, when dealers face incentives to increase sales; whether certain dealers were more likely to offer discounts in a manner correlated with customer demographics; whether the vehicle was in greater demand — measured by the average time the particular model stayed on dealers lots — and whether there was a trade-in vehicle associated with the transaction.

We uncover an important interaction between age and gender in the market for new cars. First, we find that dealer margins rise among older customers, even after controlling for all observable aspects of the transaction. Second, we find that the age premium rises more steeply for women than for men. As a result, young women pay the least for new cars whereas older women pay the most. The under- and over-payments by these groups of women are not only significantly different from zero and from each other, but also significantly different than the premiums paid by men of the same age. Thus, even if there exist secular age effects that may explain our results to some extent, the differences in outcomes between younger and older women are even more pronounced and require further explanation.

We believe that our results are consistent with the education and labor market disparities between women of different age cohorts. Women above the age of 65, who generate the highest dealer margins, are less likely to have participated in the labor force when they were younger, or to have obtained a college degree, than women under 25, who generate the lowest margins. These differences can potentially cause older women to have lower information in the new car market and possibly also to negotiate with less confidence. The price differences that we observe may be exacerbated by demographic trends. Beyond the age of 65, women start to rapidly outnumber men in the population, due to lower survival rates for the latter. Thus, a large number of women in this cohort are single, either because of the death of a spouse or because of divorce. Many of these women may have relied on their spouses to purchase new cars in the past — we see evidence that this is true even today, by the considerably lower rates of women purchasing cars among the age group most likely to be married. Thus, older women who buy new cars may be doing so for the first time, in contrast to men of the same age.

We also consider alternative explanations for our results, in particular the possibility that the higher age premium for women may be an ongoing phenomenon. The evidence supports our hypothesis that it is in fact driven by labor market and education changes, but we acknowledge that our results do not conclusively prove this. Doing so would require a time series of vehicle transactions involving consumers of different ages; essentially multiple repetitions of our data over a very long period of time. But such data do not exist in the new car market — since large samples of vehicle sales have become available only in recent years — and certainly do not exist in other markets, due to the difficulties described above. Nevertheless, ours is the first study to examine this issue and to provide evidence that the gender disparities may be caused by women's education and labor market changes. Moreover, our results are cleanly identified and extremely robust. Confirmation of this result will require it to be repeated in future studies.

In Section 2 we present the data used in our study. In Section 3 we present the empirical framework. Section 4 contains our main results along with various checks for the robustness of our findings. Section 5 discusses various explanations for our results. We provide a brief conclusion in Section 6.

2 Data

Our analysis uses data provided by automobile dealers to a major market research firm. This data includes more than 250 key observations for each vehicle transaction including: *a*) vehicle characteristics: vehicle trim, number of doors, engine type, transmission etc; *b*) transaction characteristics: suggested retail price, transacted price, cost of the vehicle including factory and dealer installed accessories, whether the vehicle was financed, leased or a cash purchase, rebates offered, and how long the vehicle was on the dealers lot; *c*) trade-in characteristics: the price of the trade-in vehicle, and its under or over valuation; *d*) customer demographics: gender, age, province of residence, and province of purchase. The selection of dealers is not random, as dealers must agree to be included in the database. However, it is the most comprehensive dataset on vehicle purchase transactions. Our sample of the data is approximately 20% of new vehicle sales in Canada from May 1st, 2004 to April 15th, 2009, a total of 1,137,573 transactions.

The automobile industry in Canada is very similar to its US counterpart. The two markets are in fact extremely integrated, and the set of manufacturers and models is almost exactly the same, as is the process of customer negotiation for new cars with dealers. Canadian prices tend to be somewhat higher than in the US, but in most other respects the two industries are indistinguishable. We have no reason to believe that our findings are unique to the Canadian autombile industry.

2.1 Sample Selection

We comparing our sample of transactions data with aggregate sales data provided by Desrosiers Automotive Consultants, and population data from Statistics Canada in order to determine how representative our sample is. Consider the distribution of transactions by major manufacturer (Table 1). The first two columns in Table 1 are the number of transactions and percentage of overall transactions for the year 2006 from our sample of the PIN data. The third column is the percentage of sales by manufacturer from data compiled by Desrosiers Automotive Consultants for the same year—2006, capturing all sales in Canada. Our sample appears to over-represent sales by GM, Toyota, Nissan, and VW. Similarly, our sample under-represents Chrysler, Ford, Honda, and other manufacturers. THe differences are, however, not large.

Next, we illustrate the distribution of vehicle transactions across provinces for the year 2006 (see Table 2). The first two columns are the number of transactions and percentage of total transactions for the PIN data. The third column is the sales share calculated from data provided by Desrosiers, and the fourth column has the percentage of population for each province estimated by Statistics Canada.[6] Our sample clearly over-represents Alberta, and to a certain extent Quebec, at the expense of slight under-representation of the other provinces.

These tables most likely illustrate dealer sample selection issues. For this reason, we are careful not to draw conclusions across provinces and manufacturers. Our estimating strategy (discussed later) focuses on effects *within* provinces and car models. Thus, sample selection

[6]Source: Statistics Canada, Canadian Socioeconomic Database, table 051-0001.

Manufacturer	No.	Column %	Aggregate Sales %
GM	53458	28.1	26.0
Toyota	31859	16.8	12.2
Chryl	22703	11.9	13.7
Ford	20934	11.0	14.2
Honda	18732	9.9	10.3
Nissan	9603	5.1	4.1
Korea	7232	3.8	6.2
Mazda	9001	4.7	5.0
VW	6689	3.5	2.7
Japan	5083	2.7	2.4
Europe	2633	1.4	1.9
BMW	2174	1.1	1.2
Total	190101	100.0	100.0

Source: Authors Calculations-JD Power PIN, and Desrosiers Data

Table 1: Sales Across Manufacturers (2006).

Province	No.	Column %	Sale-share %	Pop-share %
Ontario (ON)	65453	34.4	37.5	38.9
Quebec (QC)	51728	27.2	24.7	23.4
Alberta (AB)	42660	22.4	14.9	10.5
British Columbia (BC)	19156	10.1	11.7	13.0
Newfoundland and Labrador (NL)	2501	1.3	1.5	1.6
Manitoba (MB)	3310	1.7	2.7	3.6
Nova Scotia (NS)	1904	1.0	2.9	2.9
Saskatchewan (SK)	2040	1.1	2.4	3.0
NorthWest Territories (NT)	710	0.4	NA	0.1
New Brunswick (NB)	487	0.3	2.1	2.3
Yukon (YT)	152	0.1	NA	0.1
Total	190101	100.0	100.0	100.0

Source: Authors Calculations, PIN and Desrosiers Data, STATCAN CANSIM Table 051-0001.

Table 2: Sales Across Provinces (2006).

Segment	Mean						Median			
	Male Cust	Female Cust	Cust Price	Dealer Marg	Turn Days	Cust Age	Cust Price	Dealer Marg	Turn Days	Cust Age
Compact (37%)	46.7%	53.3%	$19,936	$1,314	44.8	44.5	$19,529	$1,238	14	44
Full (0%)	75.8%	24.2%	$35,285	$1,204	55.3	67.8	$35,215	$1,224	8	70
Luxury (4%)	68.9%	31.1%	$47,712	$2,701	53.4	50.4	$44,968	$2,550	19	50
Midsize (13%)	62.0%	38.0%	$28,617	$1,647	58.9	51.3	$27,871	$1,555	20	51
Pickup (13%)	83.0%	17.0%	$39,315	$2,209	61.0	43.4	$38,604	$1,865	26	43
SUV (20%)	59.9%	40.1%	$35,555	$1,943	49.8	45.6	$32,219	$1,702	16	45
Sports (1%)	67.6%	32.4%	$42,690	$3,075	58.9	43.1	$35,041	$2,240	18	44
Van (9%)	67.1%	32.9%	$28,423	$1,346	71.3	47.8	$27,800	$1,361	36	45
Total (100%)	59.6%	40.4%	$29,255	$1,705	53.0	46.0	$26,373	$1,463	18	46

Source: Authors Calculations

Table 3: Statistics across Segments.

should not be a concern for the implications of our results.

In Table 3 we present selected summary statistics across vehicle segments. The bottom row presents summary statistics for the entire dataset. The median vehicle sold for approximately $26,373, generated $1463 in margin for the dealer, was on her lot for 18 days, and was bought by a 46 year old male customer. Full size (with relatively few data-points) and luxury vehicles are purchased by older individuals. Luxury vehicles have the highest values. Compact and full size vehicles are sold in fewer days than other segments.

In the next few tables we highlight the role played by gender and age in the automobile market. In our data, gender is determined by matching the entered name with expected gender. There are several observations where gender or age cannot be determined, and we drop these. We also drop observations with very small and very large dealer margins; in particular the bottom and top 1% of the data. These correspond to unusual models, typically with very high purchase prices, which can potentially skew the results.[7] This leaves us with about 510,000 observations (from over a million earlier). All summary statistics presented in this data description, as well as our later regressions, are from this smaller subset.

Approximately, 59% of all transactions have male customers (see Table 4). We divide age into 11 categories of 5-year intervals. Approximately 24% of our transactions have customers under the age of 35. Women outnumber men in the under 25 age category, but not in any other category. We observe an interesting 'missing women' pattern. Women comprise 46% of transactions among the under 35 year-olds. This shrinks to 41% in midage category, and further to 36% in the category aged above 55. With little other information, it is difficult to hypothesize an explanation.

In Figure 1 we illustrate how the mean transacted price, and vehicle cost varies with age in our data. Both mean price and cost peak for customers aged 39 years, at $30,306, and $28,900 respectively. Both younger and older customers around this age buy cheaper

[7]This procedure drops transactions where the dealer margin was greater than $7300, or where the dealer lost more than $2200 on the sale.

Age Category	Gender								
	Male			Female			Total		
	No.	Col %	Row %	No.	Col %	Row %	No.	Col %	Row %
Age Under 25	15301	5.1	47.8	16716	7.9	52.2	32017	6.3	100.0
Age 25-30	22801	7.6	52.3	20776	9.8	47.7	43577	8.5	100.0
Age 30-35	28251	9.5	58.0	20493	9.6	42.0	48744	9.5	100.0
Age 35-40	31889	10.7	60.0	21270	10.0	40.0	53159	10.4	100.0
Age 40-45	36553	12.2	58.1	26379	12.4	41.9	62932	12.3	100.0
Age 45-50	38623	12.9	57.4	28631	13.5	42.6	67254	13.2	100.0
Age 50-55	35561	11.9	58.0	25789	12.1	42.0	61350	12.0	100.0
Age 55-60	29708	10.0	59.8	19976	9.4	40.2	49684	9.7	100.0
Age 60-65	21921	7.3	62.4	13231	6.2	37.6	35152	6.9	100.0
Age 65-70	15341	5.1	64.3	8531	4.0	35.7	23872	4.7	100.0
Age Over 70	22554	7.6	68.1	10571	5.0	31.9	33125	6.5	100.0
Total	298503	100.0	58.4	212363	100.0	41.6	510866	100.0	100.0

Source: Authors Calculations

Table 4: Sales by Age Category

vehicles on average.

A similar trend in dealer margin is illustrated in Figure 2. Like the average price and cost, dealer margin rises with age and then declines after a peak of $1769 at age 50. This distribution is marked by a relatively flat section from age 33 (average dealer margin $1707) to age 59 (average dealer margin $1707) with sharper drop offs on either side.

Overall, men purchase more expensive vehicles, with an average price of $31,075 versus $26,589 for females (see Table 5). Male customers also generate higher margins for their dealers at $1,784 versus $1,590 generated by female customers. This difference carries through within vehicle segments as well.

We can also illustrate how vehicle prices and cost vary across age by gender. In Figure 3 we illustrate how the mean transacted price, and vehicle cost varies with age across gender. Male customers aged 42 years, buy the most expensive vehicles at $31,709, and female customers aged 39 years buy their most expensive vehicles at $26,781. Both younger and older customers around this age buy cheaper vehicles on average.

3 Estimating Strategy

Consider a simple linear model explaining dealer profits from a particular vehicle transaction. Define x_m as model characteristics, x_d as dealer characteristics, and z_g as a group identifier $\forall g \in 1, 2$ (for simplicity we assume only two groups). Dealer d's profits from selling model m to customer i from group g are:

$$\pi^m{}_{idg} = x_m\beta + x_d\gamma + z_g\delta + \epsilon_i, \tag{1}$$

Figure 1: Price and Vehicle Cost by Age

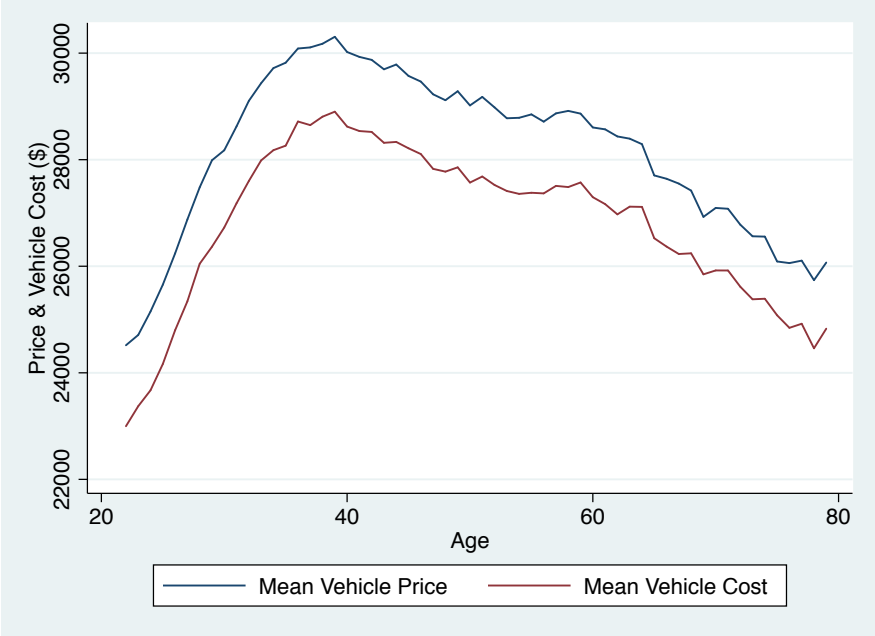

where β, γ, δ are the associated coefficients, and ϵ_i is the error term. We identify a market disparity across groups if $\delta \neq 0$.

3.1 Omitted Variable and Misspecification Bias in Estimating Discrimination.

Equation 1 above helps us identify a market disparity, it does not help us identify 'discrimination.' Consider a special case for the two groups, where the true value of $\delta = 0$ and $\beta_1 = \beta + \phi$, while $\beta_2 = \beta$. In other words, due to a difference in preferences, group one assigns model characteristics a different value than group two. If we estimated equation 1 naively we would mistakenly estimate $\hat{\delta} = x_m\phi$, assigning discrimination where only a difference in preferences explains the market disparity.

3.2 Empirics

Our empirical strategy will be a simple reduced form. Following Busse et al. (2013), a vehicle price equation is estimated by regressing it on demand covariates (X_D) and supply covariates (X_S),

$$P = \alpha_0 + X_D\alpha_1 + X_S\alpha_2 + \nu. \tag{2}$$

Demand covariates include demographic variables (what we are most interested in). Other demand and supply covariates are controlled for by a series of fixed effects, explained below.

Our empirical strategy, relies on the richness of our data. We will not fit a structural model to estimate price elasticities and optimal markups. We only aim to robustly identify variations as present in our data.

Figure 2: Dealer Margin by Age

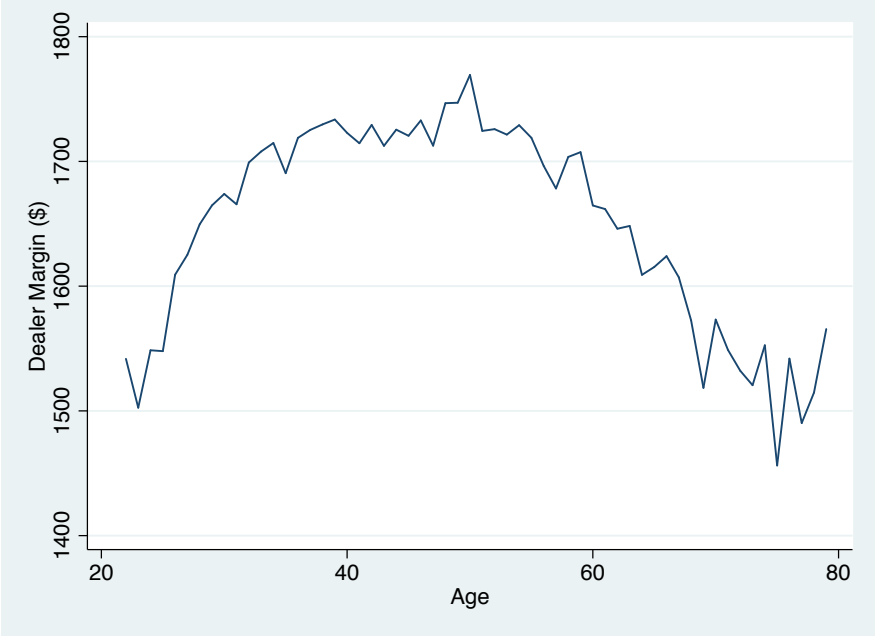

Segment	Mean			
	Cust. Price	Dealer Marg	Cust. Price	Dealer Marg
Compact (37%)	$20,183	$1,327	$19,720	$1,302
Full (0%)	$35,342	$1,217	$35,108	$1,164
Luxury (4%)	$48,590	$2,740	$45,773	$2,614
Midsize (13%)	$28,844	$1,650	$28,247	$1,643
Pickup (13%)	$39,446	$2,202	$38,678	$2,240
SUV (20%)	$36,263	$1,957	$34,496	$1,922
Sports (1%)	$45,141	$3,335	$37,627	$2,542
Van (9%)	$28,446	$1,337	$28,377	$1,365
Total (100%)	$31,075	$1,784	$26,589	$1,590

Source: Authors' Calculations

Table 5: Mean Customer Price and Dealer Margin by Gender.

Figure 3: Price and Vehicle Cost by Age and Gender

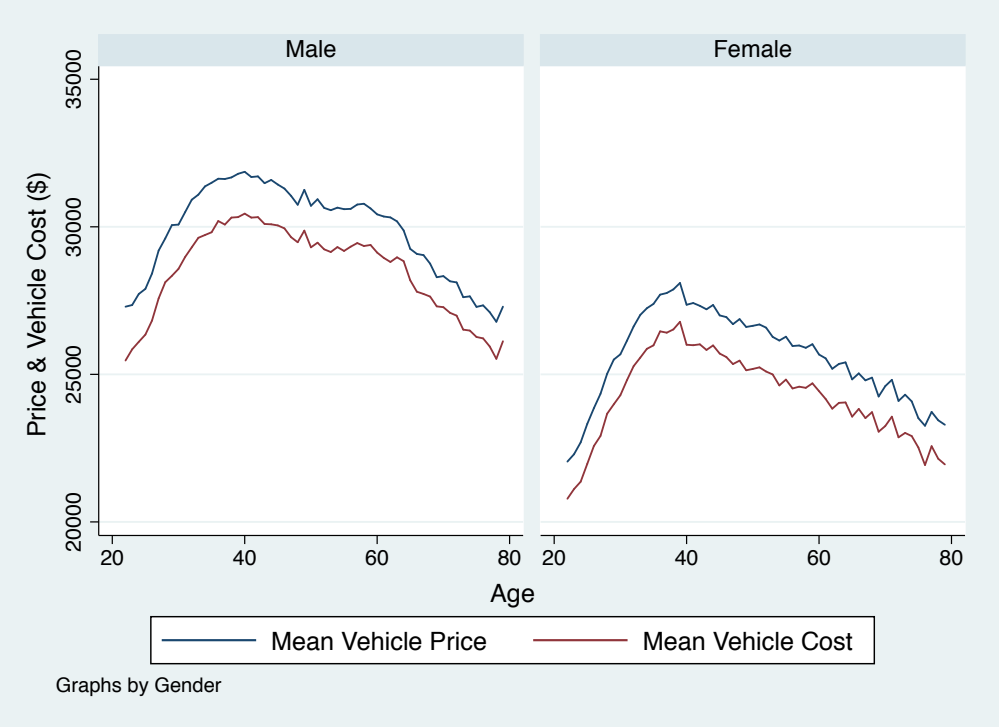

For our reduced form analysis we believe two considerations are important. The first is to identify variations in price paid and profit generated within as uniform a product as possible. Our very large sample size is an incredible asset for this. Initial analysis of the data reveals that there are many high selling vehicle models, with over 1,000 transactions for each model, year combination (over 5,000 transactions overall). These popular models (such as the Toyota Corolla and Honda Civic) draw customers from a wide distribution of gender and age. Our data will allow us to robustly identify the within model, year, province and trim variations in vehicle price and dealer profits across several gender and age categories. Our second consideration is to minimize the potential bias from omitted variables. Recent analysis of the US version of the PIN indicates that besides model and dealer characteristics, transaction characteristics influence price and dealer profits. Whether the vehicle was leased, or financed, or if the vehicle was bought at the end of the month, or year, and whether there was a trade-in vehicle associated with the transaction can have important effects. If a gender and age category displays a substantially distribution over these transaction characteristics than the rest of the population, omitting them would bias our results. Such variables, or suitable proxies, will be included as controls.

Our estimating equation is an adaptation of the reduced form used in Busse et al. (2013). Dealer profits (and equivalently price) can be estimated as:

$$\pi^{m}_{idgpt} = \beta_0 + \beta_1 \text{demog}_{it} + \beta_2 \text{timing}_{it} + \beta_3 \text{transchar}_{it} + \tau_m + \mu_d + \lambda_t + \epsilon_{it} \qquad (3)$$

where $\beta_0, \beta_1, \beta_2, \beta_3$ are the associated coefficients, and ϵ_{it} is the error term for individual i at time t. Our coefficient of interest would be β_1 identifying a market disparity based on the demographic group indicator. Included in this equation will be model, province, year and

Figure 4: Customer Age Distribution for Popular Models

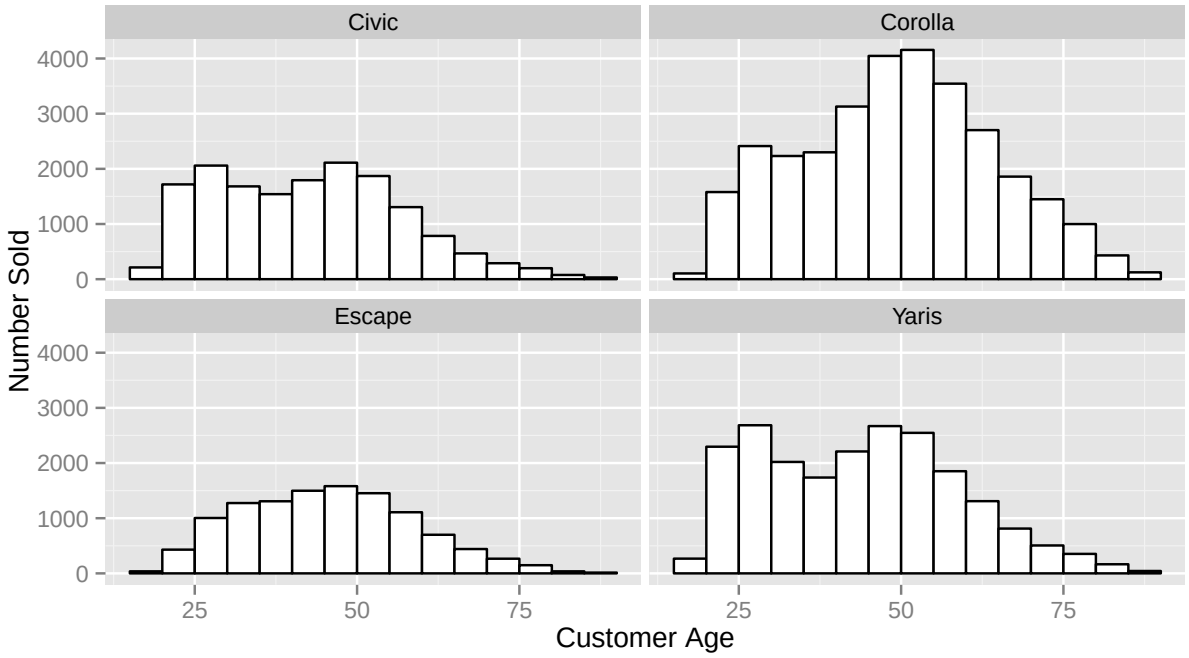

trim fixed effects τ_m, dealer fixed effects μ_d and time fixed effects λ_t. This equation will yield the average within model-province-year-trim difference in dealer margin across our gender and age categories.

For illustration, assume that in our dataset we have two models: a Honda Civic, and a Toyota Corolla, two provinces: Ontario, and Quebec, two years: 2008 and 2009, and finally two demographic categories: male and female. The average difference in dealer margin across males and females for a Honda Civic sold in Ontario in 2008 would be a part of our estimate. The estimate generated by our equation, further averages this average difference across males and females over each of the 8 possibilities in the combinations above.

3.3 The reasons for market disparities across gender and age.

On discovering market disparities across gender and age groups how are we to interpret them. Are differences in dealer prices due to statistical or taste based discrimination? The possibilities of statistical discrimination are several. A variation in price elasticity, and preferences across gender and age could explain the difference in market outcomes observed. Given the importance of negotiation in the market, an inherent difference in bargaining strategies across gender (see Babcock and Laschever (2003)) could yield differences. Differences in search costs—through internet access, or a propensity to research vehicles—across gender and age could influence market outcomes (similar to Morton et al. (2003)).

A taste based explanation for differences in market outcomes derives from differences in the utility derived from interacting with different groups. For instance, a male dominated car salesman might prefer interacting with a young female customer and be willing to accept a smaller profit from that transaction.

Table 6: Regression of Log(Vehicle Cost)

	(1)		(2)		(3)	
Female	-0.123^a	(0.001)	-0.000	(0.000)	0.003^a	(0.000)
Age 25-30	0.091^a	(0.002)	0.003^a	(0.001)	0.003^a	(0.001)
Age 30-35	0.162^a	(0.002)	0.005^a	(0.001)	0.006^a	(0.001)
Age 35-40	0.201^a	(0.002)	0.007^a	(0.001)	0.008^a	(0.001)
Age 40-45	0.202^a	(0.002)	0.007^a	(0.001)	0.008^a	(0.001)
Age 45-50	0.188^a	(0.002)	0.008^a	(0.001)	0.008^a	(0.001)
Age 50-55	0.180^a	(0.002)	0.009^a	(0.001)	0.009^a	(0.001)
Age 55-60	0.184^a	(0.002)	0.010^a	(0.001)	0.010^a	(0.001)
Age 60-65	0.183^a	(0.003)	0.009^a	(0.001)	0.011^a	(0.001)
Age 65-70	0.160^a	(0.003)	0.005^a	(0.001)	0.010^a	(0.001)
Age $>$ 70	0.117^a	(0.003)	-0.000	(0.001)	0.008^a	(0.001)
Financed Indicator	-0.033^a	(0.001)	0.019^a	(0.001)	0.022^a	(0.000)
Leased Indicator	0.031^a	(0.001)	0.027^a	(0.001)	0.029^a	(0.000)
Sat or Sun FE	-0.005^a	(0.001)	-0.001^a	(0.001)	-0.001^b	(0.000)
End of Month FE	0.015^a	(0.001)	-0.001^a	(0.000)	-0.001^a	(0.000)
End of Year FE	0.010^b	(0.004)	-0.001	(0.002)	-0.001	(0.001)
Constant	9.987^a	(0.007)	10.223^a	(0.003)	10.199^a	(0.002)
R^2	0.171		0.896		0.935	

^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$. Fixed-effects specified as Col. 1: Province; Col. 2: Model-Year-Province; Col. 3: Model-Year-Trim-Province. All regressions include year-month FEs. N=510866.

A combination of several factors likely explains observed market disparities in the new automobile market. It is unlikely that we will have a certain answer on the relative contribution of different explanations. While our research will look for associated evidence for various explanations, that will not be its primary goal.

4 Regression Results

We now present results from estimating Equations XX and XXX in the previous section.

4.1 Regressions of Vehicle Cost

We first examine how the age and gender of consumers relates to vehicle costs. Table 6 regresses the log of the vehicle cost, as recorded by the dealer, on demographic characteristics. Column 1 includes fixed-effects for the province of the transaction. Column 2 uses Model*model-year*province fixed-effects, and Column 3 uses Model*model-year*trim*province fixed-effects. All regressions include year-month fixed-effects to control for general trends in vehicle prices. The first column reveals the relationship between age and gender and average new car prices, *without* controlling for the make or model or trim-level of the car. The omitted age category is consumers aged below 25. Thus the results in column 1 show that the

youngest consumers buy the cheapest cars. Consumers aged 25-30 buy 10% more expensive cars on average, and the premium is about 18–20% for the other age groups. Consumers aged 70 and above buy less expensive cars than consumers in all age categories between 30 and 70.

Column 2 of Table 6 shows that, once we control for the vehicle model, there are no longer large differences across age groups. While the very youngest and oldest consumers purchase less expensive trims, the difference between these consumers and the rest is at most one percent. A similar pattern in column 3 shows that there is no clear relationship between customer age and aftermarket options or warranties, other than evidence that consumers under the age of 30 are the least likely to make such purchases.

Column 1 shows that female consumers buy new cars that are 12% cheaper on average than those bought by male consumers. However, this difference completely disappears once we control for model in Column 2, indicating that gender has no relationship with vehicle trim within a given model. Column 3 suggests that female consumers spend slightly more on options and warranties than men do, although the difference of 0.3% is economically very small despite being statistically significant.

We also include indicators for the timing of the transaction, and whether the vehicle was leased or financed through the dealer. Column 1 shows that consumers who finance their vehicles tend to buy cheaper cars. However, once we control for make in column 2 and trim in column 3, consumers who finance tend to buy more expensive vehicles. A similar effect is seen for consumers who lease vehicles. Weekends see the sales of slightly cheaper vehicles on average, although once make and trim are controlled for, the effect is tiny. On average, slightly more expensive vehicles are sold at the end of the month and the year, but again the effects are small or non-existent when we control for model and trim.

In Table 7 we examine the joint distribution of age and gender, using the same sets of fixed-effects as in the previous table. The omitted category is male consumers under the age of 25. The results in column 1 show that in every age category, female consumers buy cheaper cars on average than their male counterparts of the same age, showing that the previous result regarding gender holds across age groups. Indeed the male-female new car premium is a remarkably robust 13–16% in almost every age category. Young women buy the cheapest cars of any gender-age group, while 35 to 65 year old men buy the most expensive cars.

The large coefficients that we observe in column 1 of Table 7 disappear in column 2, which controls for the model purchased. Thus, there appears to be no systematic relationship between customer demographics and the choice of trims within a given model, just as we saw in Table 6. In column 3, the coefficients on each demographic group are positive and statistically significant, indicating that men under the age of 25 spend the least on aftermarket options and warranties. Examining the coefficients suggests that, beyond the age of 40, women spend slightly more on these add-ons than men of the same age, but the differences are small.

Our results so far indicate that consumer demographics are systematically related to the purchase price of the vehicle, but that this is almost entirely due to choices regarding which make and model to purchase. Conditional on the model, demographics do not have a clear relationship with consumers' choice of vehicle trims or other options. We now turn to the question of how consumers' characteristics relate to dealer profits.

Table 7: Regression of Log(Vehicle Cost)

	(1)		(2)		(3)	
Age < 25 Female	-0.157^a	(0.004)	-0.007^a	(0.001)	0.003^a	(0.001)
Age 25-30 Male	0.077^a	(0.003)	0.001	(0.001)	0.005^a	(0.001)
Age 25-30 Female	-0.054^a	(0.003)	-0.004^a	(0.001)	0.005^a	(0.001)
Age 30-35 Male	0.140^a	(0.003)	0.003^a	(0.001)	0.008^a	(0.001)
Age 30-35 Female	0.027^a	(0.004)	-0.002^c	(0.001)	0.006^a	(0.001)
Age 35-40 Male	0.174^a	(0.003)	0.004^a	(0.001)	0.009^a	(0.001)
Age 35-40 Female	0.074^a	(0.003)	0.003^b	(0.001)	0.010^a	(0.001)
Age 40-45 Male	0.182^a	(0.003)	0.004^a	(0.001)	0.009^a	(0.001)
Age 40-45 Female	0.065^a	(0.003)	0.003^b	(0.001)	0.010^a	(0.001)
Age 45-50 Male	0.169^a	(0.003)	0.004^a	(0.001)	0.008^a	(0.001)
Age 45-50 Female	0.049^a	(0.003)	0.005^a	(0.001)	0.011^a	(0.001)
Age 50-55 Male	0.160^a	(0.003)	0.003^b	(0.001)	0.008^a	(0.001)
Age 50-55 Female	0.043^a	(0.003)	0.007^a	(0.001)	0.014^a	(0.001)
Age 55-60 Male	0.169^a	(0.003)	0.005^a	(0.001)	0.009^a	(0.001)
Age 55-60 Female	0.039^a	(0.004)	0.007^a	(0.001)	0.015^a	(0.001)
Age 60-65 Male	0.173^a	(0.003)	0.005^a	(0.001)	0.010^a	(0.001)
Age 60-65 Female	0.029^a	(0.004)	0.006^a	(0.001)	0.015^a	(0.001)
Age 65-70 Male	0.146^a	(0.004)	-0.001	(0.001)	0.008^a	(0.001)
Age 65-70 Female	0.014^a	(0.004)	0.004^b	(0.002)	0.016^a	(0.001)
Age > 70 Male	0.099^a	(0.003)	-0.006^a	(0.001)	0.007^a	(0.001)
Age > 70 Female	-0.025^a	(0.004)	-0.001	(0.002)	0.015^a	(0.001)
Financed Indicator	-0.034^a	(0.001)	0.019^a	(0.001)	0.022^a	(0.000)
Leased Indicator	0.031^a	(0.001)	0.027^a	(0.001)	0.029^a	(0.000)
Sat or Sun FE	-0.005^a	(0.001)	-0.001^a	(0.001)	-0.001^b	(0.000)
End of Month FE	0.015^a	(0.001)	-0.001^a	(0.000)	-0.001^a	(0.000)
End of Year FE	0.010^b	(0.004)	-0.001	(0.002)	-0.001	(0.001)
Constant	10.005^a	(0.008)	10.226^a	(0.003)	10.199^a	(0.002)
R^2	0.171		0.896		0.935	

^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$. Fixed-effects specified as Col. 1: Province; Col. 2: Model-Year-Province; Col. 3: Model-Year-Trim-Province. All regressions include year-month FEs. N=510866.

Table 8: Regression of Dealer Margin

	(1)		(2)		(3)	
Female	-107.6^a	(3.9)	13.0^a	(3.6)	21.5^a	(3.6)
Age 25-30	107.5^a	(10.0)	6.5	(9.1)	8.4	(9.0)
Age 30-35	178.0^a	(9.8)	6.7	(9.0)	8.5	(9.0)
Age 35-40	224.6^a	(9.6)	23.2^a	(8.9)	26.5^a	(8.9)
Age 40-45	254.7^a	(9.3)	45.1^a	(8.6)	48.1^a	(8.6)
Age 45-50	269.5^a	(9.2)	68.6^a	(8.5)	71.8^a	(8.5)
Age 50-55	281.2^a	(9.4)	84.8^a	(8.6)	87.4^a	(8.6)
Age 55-60	282.7^a	(9.8)	84.5^a	(9.0)	88.9^a	(9.0)
Age 60-65	257.4^a	(10.6)	79.5^a	(9.7)	83.9^a	(9.7)
Age 65-70	232.3^a	(11.7)	98.6^a	(10.8)	105.2^a	(10.8)
Age $>$ 70	194.0^a	(10.8)	130.1^a	(10.1)	144.4^a	(10.1)
Financed Indicator	-44.0^a	(6.1)	171.5^a	(5.8)	180.4^a	(5.8)
Leased Indicator	36.3^a	(6.0)	185.5^a	(5.7)	198.0^a	(5.7)
Sat or Sun FE	20.4^a	(6.1)	22.3^a	(5.5)	23.1^a	(5.5)
End of Month FE	-40.5^a	(4.4)	-60.0^a	(4.0)	-61.0^a	(4.0)
End of Year FE	-88.1^a	(17.9)	-93.2^a	(16.2)	-93.6^a	(16.2)
Constant	1607.8^a	(31.0)	2384.5^a	(26.7)	2360.3^a	(26.9)
R^2	0.037		0.239		0.264	

^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$. Fixed-effects specified as Col. 1: Province; Col. 2: Model-Year-Province; Col. 3: Model-Year-Trim-Province. All regressions include year-month FEs. N=510866.

4.2 Regressions of Dealer Margin

The results are presented in Table 8, where the dependent variable is the dealer's margin on each transaction. Note that we do not take logs of this variable because dealer margins are negative in about 8% of cases, which would require us to drop these observations.[8] Column 1 shows that dealers make on average $107 less on female consumers, as well as lower margins on younger consumers. A 45 to 65 year old consumer generates about $280 higher dealer margins than a consumer under 25, and about $100 more than those between 25 and 30. But this is not surprising, given the same pattern with regard to the kinds of cars these consumers buy, as shown in Table 6. That is, dealers make their lowest margins on the cheapest cars. More interesting are the results in the other two columns. Column 2, which controls for model, reveals that dealers generate slightly higher margins on female consumers and considerably higher margins on older consumers — as much as $100–130 on consumers above the age of 65. There differences persist, and are in fact somewhat larger, in column 3, which also controls for vehicle trim. Consumers above 65 generate about $105–144 higher margins for dealers than consumers under the age of 35.

[8] As discussed above, dealers make negative margins for a number of reasons, including responding to manufacturer incentives, getting rid of cars with low demand, and taking losses on the new car in order to make profits on the trade-in. Below, we show that our results are fully robust to using only transactions with positive dealer margins, and expressing the dependent variable in logs.

These results are interesting and perhaps surprising. We have controlled for the model, province and trim of the vehicle, and note again that the relevant fixed effects are defined as the product of these categories. In other words, we examine differences in dealer margins within the same model-trim and province, and separately also controlled for the timing of the purchase through fixed-effects for year-month, and for weekends, month-ends and year-ends. Therefore, we are examining almost exactly the same car being sold at the same time to consumers of different types, as reflected by their demographic characteristics. Thus, any differences in the margins that dealers generate in these transactions must be due to these demographics, whether because of different behavior exhibited by these groups or different dealer reactions to encountering consumers of these types.

We break down the results more finely in Table 9, where we use the joint distribution of age and gender categories. The omitted category is male customers under the age of 25. The results are very clear: dealers make their lowest margins on young, female consumers. This is most obvious in column 1, where women under 25 generate almost $300 lower margins than men of the same age, and about $450 lower margins than men aged 50–60. But this is partly driven by the choice of cars that these types of consumers purchase. When we control for the model in the transaction, women under 25 pay about $140 less than the omitted group, which is men of the same age. More interestingly, now the group generating the greatest profit for dealers is women above the age of 70, who pay an average of $85 above the omitted group.

In column 3 we control for model trim and see the same pattern of results. The lowest profits for dealers come from women under 25, followed by those between 25 and 30. The highest profits are generated on women above 70, followed by those in the 60–65 age group.

In general, the results show that dealer margins rise with the age of consumers, as we see that older men also pay more than average. However, the striking result is that, while young women do better than men of the same age, older women do considerably worse. Thus, not only does the premium paid by consumers rise with age, but it does so much more steeply for women than for men. We show this graphically in Figure 5.

Overall, dealer margins are over $100 higher on older female customers than on the omitted category of young men, and about $115 lower for young women than these men, using the coefficients in Column 3. While these amounts may possibly appear small on the purchase of a car that costs upwards of $20,000 it is worth noting two points. First, these simply reflect the realistic amounts that most buyers can hope to save by conducting research, obtaining multiple quotes and negotiating accordingly. It is unlikely that buyers can expect to save much more than a few hundred dollars on a given purchase. Second, the amounts we estimate are averages for each demographic group, implying that some consumers may well enjoy considerably larger savings through better strategies.

Turning to the other factors that may influence dealer margins, we see from the third column in both Tables 8 and 9 that dealers make $180–200 higher margins on cars that are leased or financed through them, relative to outright cash purchases. This accords with common observations about the new car market. Finally, dealers clearly sacrifice margins at the end of the month and year. On average, their margins are $61 and $94 lower at these times, again conforming to casual observations, and confirming that dealers respond to manufacturer incentives at the end of calendar months and years. Note that these lower margins may either be due to dealers being willing to sacrifice profits to to meet sales quotas,

Table 9: Regression of Dealer Margin

	(1)		(2)		(3)	
Age < 25 Female	-291.6^a	(15.2)	-138.4^a	(13.8)	-115.3^a	(13.7)
Age 25-30 Male	58.2^a	(14.2)	-34.7^a	(12.8)	-26.7^b	(12.8)
Age 25-30 Female	-148.0^a	(14.4)	-101.4^a	(13.1)	-82.1^a	(13.1)
Age 30-35 Male	92.0^a	(13.6)	-58.2^a	(12.4)	-49.3^a	(12.4)
Age 30-35 Female	-40.1^a	(14.5)	-79.7^a	(13.2)	-61.8^a	(13.2)
Age 35-40 Male	127.4^a	(13.3)	-54.7^a	(12.2)	-44.1^a	(12.2)
Age 35-40 Female	22.1	(14.4)	-45.9^a	(13.1)	-26.7^b	(13.1)
Age 40-45 Male	146.9^a	(13.1)	-43.8^a	(12.0)	-31.8^a	(12.0)
Age 40-45 Female	66.5^a	(13.8)	-8.9	(12.6)	7.8	(12.6)
Age 45-50 Male	159.8^a	(13.0)	-25.2^b	(11.8)	-13.5	(11.9)
Age 45-50 Female	83.6^a	(13.6)	21.1^c	(12.4)	38.5^a	(12.4)
Age 50-55 Male	168.0^a	(13.1)	-6.2	(12.0)	4.3	(12.0)
Age 50-55 Female	100.6^a	(13.9)	34.0^a	(12.7)	51.5^a	(12.7)
Age 55-60 Male	166.7^a	(13.5)	-15.6	(12.4)	-2.3	(12.4)
Age 55-60 Female	108.1^a	(14.6)	48.0^a	(13.3)	65.9^a	(13.3)
Age 60-65 Male	154.2^a	(14.3)	-13.6	(13.1)	-3.3	(13.1)
Age 60-65 Female	65.0^a	(16.2)	34.9^b	(14.7)	57.2^a	(14.7)
Age 65-70 Male	128.2^a	(15.6)	13.3	(14.2)	25.9^c	(14.2)
Age 65-70 Female	42.4^b	(18.4)	42.0^b	(16.7)	66.2^a	(16.7)
Age > 70 Male	79.2^a	(14.3)	40.1^a	(13.2)	63.0^a	(13.3)
Age > 70 Female	29.5^c	(17.3)	85.0^a	(15.8)	112.2^a	(15.8)
Financed Indicator	-43.6^a	(6.1)	172.0^a	(5.8)	180.9^a	(5.8)
Leased Indicator	37.2^a	(6.0)	186.3^a	(5.7)	198.7^a	(5.7)
Sat or Sun FE	20.2^a	(6.1)	22.3^a	(5.5)	23.1^a	(5.5)
End of Month FE	-40.4^a	(4.4)	-59.8^a	(4.0)	-60.9^a	(4.0)
End of Year FE	-87.9^a	(17.9)	-93.1^a	(16.2)	-93.6^a	(16.2)
Constant	1703.5^a	(31.9)	2463.4^a	(27.5)	2432.0^a	(27.7)
R^2	0.038		0.240		0.264	

^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$. Fixed-effects specified as Col. 1: Province; Col. 2: Model-Year-Province; Col. 3: Model-Year-Trim-Province. All regressions include year-month FEs. N=510866.

Figure 5: Dealer Margin premium for demographic groups, relative to Men under 25

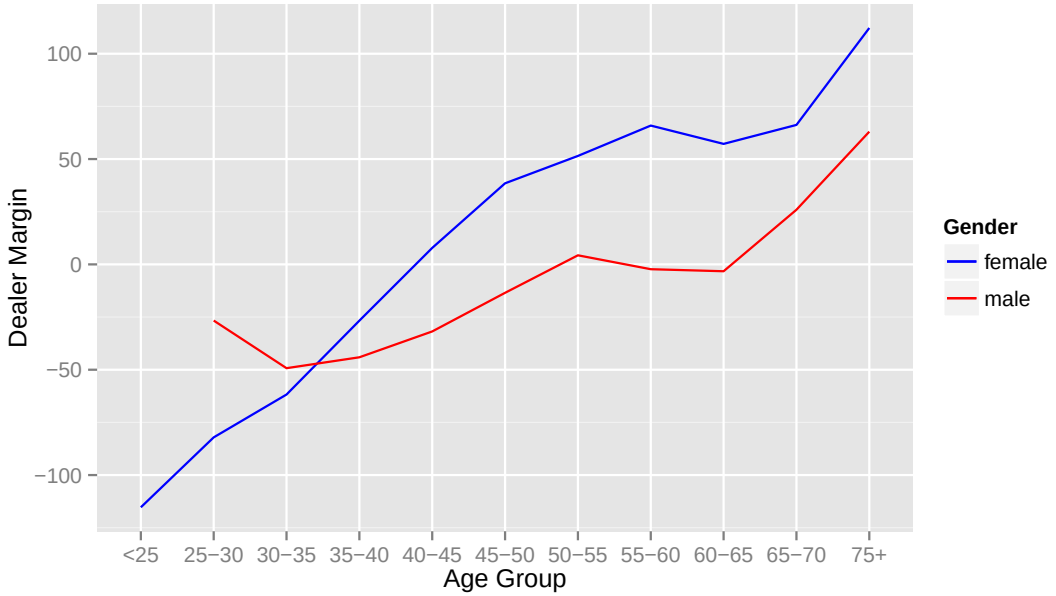

or by more price-conscious consumers negotiating transactions at these times, knowing the incentives dealers face.

4.3 Robustness Regressions

Below, we will discuss the possible reasons for our results on particular demographic groups. First, though, we present evidence that the results are not driven by other factors, such as omitted variables, outlier observations, or selection. The regression results for the following exercises are presented in Tables 10, 11 and 12 in the Appendix. We present the main coefficients of interest more compactly in Figure 6. Note that these results are obtained from regressions using the same specification as Column 3 of Table 9, which include fixed effects for model-year-province-trim and for year-months.

We start by examining particular segments — Compact cars, SUVs and Vans in Panels A and B. These segments are disproportionately associated with particular age groups, as we saw in Section 2; younger consumers are more likely to purchase small cars, while middle-aged consumers, especially those that are married, are more likely to purchase family cars such as SUVs and minivans. We then show, in Panel C, that our results are robust to examining only popular car models; i.e. that they are not driven by unusual consumer or dealer behavior in vehicles with low sales. We define high-selling car models as those with at least 5000 units in sales during our sample period. This corresponds to 26 models, out of our observed set of more than 500 models, and these comprise almost 50% of total vehicle sales.

Next, we consider the possibility that new vehicle transaction prices may be influenced by the amounts negotiated on consumers' trade-in vehicles. While in principle the two transactions should be treated separately, in practice dealers may allow overpayment on

Figure 6: Dealer Margin premium for subsamples, relative to Men under 25

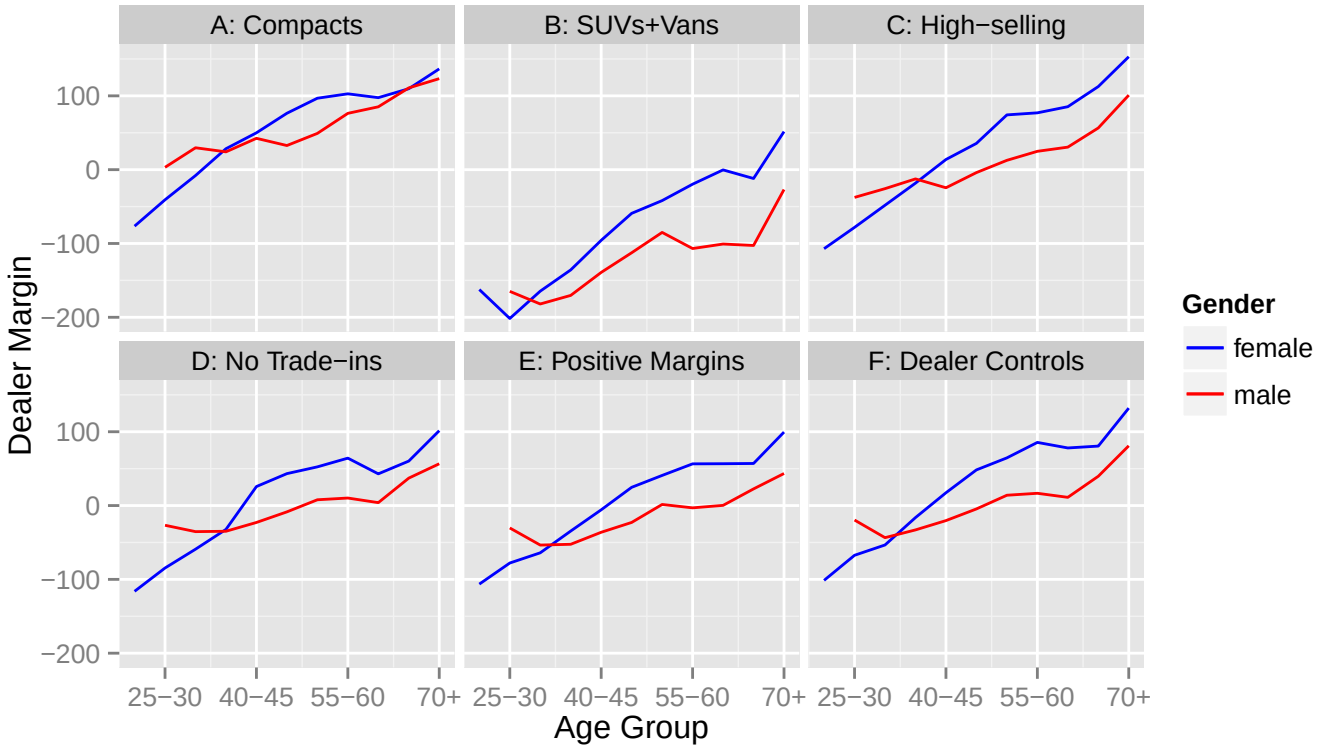

some trade-ins in order to make greater margins on the new car, or vice versa.[9] We therefore restrict attention, in Panel D of Figure 6, to the subset of transactions where consumers did not trade-in an older vehicle. We then restrict the sample, in Panel E, to those where dealers make positive margins on new car sales. While we have explained above that there may be rational reasons for dealers to make losses on certain transactions, one may be concerned that these sales are specific to certain types of cars or to unobserved characteristics of the transaction which may be correlated with consumer demographics.

Finally, we control for systematic pricing behavior by certain dealers. Suppose, for example, that dealers located in smaller cities or suburban locations charge lower prices for exactly the same car as a city dealer who faces higher costs. If the demographic distribution of consumers in these two locations is also different — for example, if older consumers are more likely to live near high-cost dealers — then this selection of consumer types may drive our results. We can control for this possibility by including dealer fixed-effects, which allow us to look *within* dealers. We also include the number of days that the vehicle has been on the dealer's lot — popular cars typically turnover very quickly and so dealers may be willing to reduce margins on cars that are not in high demand. The coefficients are presented in Panel F of Figure 6.

All six panels support our overall finding that dealer margins rise with consumer age; this is a very robust result. With regard to our second main result — that the age premium rises more steeply for women than for men — we also find considerable support. In all

[9] See Zhu et al. (2008) for evidence regarding this possibility.

six panels, the blue line corresponding to the female premium starts below the red line corresponding to the male premium, and ends up above it. However, this result is less robust in some subsamples. The gender gap among older consumers is quite small when we examine compacts alone in Panel A. In addition, the coeffcients for older women are not always significantly different from zero in Panel B (see Column 2 of Table 10 for detailed results and standard errors). Generally, though, the finding regarding differential age effects by gender is confirmed.

5 Discussion of results

In this section we discuss some possible explanations for the regression results obtained in Section 4. Recall our two main findings: first, dealers generally make higher margins on older consumers. Second, the age premium rises more steeply for women than for men; as a result, young women pay the least for new cars, whereas older women pay the most.

We first discuss the finding of a secular correlation between customer age and dealer margin.[10] Recall that our results are obtained from the period 2004–2009, and so essentially represent a snapshot of the new car market at a particular point of time. Therefore, we cannot say with certainty whether this result will hold continuously, or is particular to the time period that we examine. In other words, the result may either represent an *ongoing* effect, in which case we should expect older consumers to continuously fare worse in the new car market, or it may be a *cohort* effect, which is particular to the set of consumers that were observed during this period.

If the ongoing effect is true — i.e. if older consumers simply do worse in price negotiations — then we would expect this result to be repeated in future studies as well. In particular, we would expect the young consumers of today, who currently negotiate the lowest prices, to do worse as they age. By contrast, if the cohort effect is true, then there is something particular about the current generation of older consumers that leads them to fare worse in the new car market, but this fact may not continue to hold for future generations.

Both explanations may have some support. We know from prior research (Morton et al. (2011)) that search costs are important in price negotiations for new cars. It is possible that younger consumers have the lowest search costs and that search costs rise with the age of consumers. However, it is perhaps a little counter-intuitive to think of consumers above 65 as having the highest search costs, as these consumers are most likely to be retired or out of the labor force.

Prior research has also shown that research on the Internet can reduce the prices that consumers pay for new cars. (Morton et al. (2001)) show that consumers using an online referral service can do better than a majority of other users. Since that study, there have been a host of new sources of information on new car sales. Websites such as Carfax.com and Edmunds.com provide consumers with detailed information about the invoice price of particular vehicles, suggesting that the savings from online research may have grown considerably. Separately, there is clear evidence that older consumers are less likely to use

[10]Note that a similar finding was obtained by Harless and Hoffer (2002), although they did not break down their result by gender. Moreover, that study used a much smaller sample than ours, and did not control for other factors such as idiosyncratic dealer behavior, trade-in vehicles, etc.

the Internet (cite StatsCan data). If this is true then the age effects we estimate may reflect existing gaps in online use between old and young consumers, but these may not persist as the current cohort of young users — who are comfortable going online — grows older.

We now turn to the differential age effects by gender. Ours is the first study to identify this effect and as with the first finding there is no clear explanation. However, we believe it is unlikely that this result is an ongoing effect: there is no good reason to expect that the young women of today, who are negotiating better autombile prices than men of the same age, will in the future start to do worse than these same men. In other words, if we treat male consumers as a control group, then the results are evidence that there is something different about young women in our sample compared to older women.

Our explanation builds on the very different labor market experiences and education attainment of women in the under 30 cohort versus those in their 60s and above. It is well known that over the past few decades, women in North America have considerably narrowed the education gap with men, and in recent years have even surpassed men. At the same time, women's labor force participation has grown rapidly.

The economic effects of these societal changes are important, and potentially explain our findings. Most prior research on these economic consequences has focused on labor market outcomes — most importantly, on women's wages. Research has shown that while there has been some reduction in the pay gap, women still earn less than men in similar occupations. There is considerably less research on the impacts this has had on other economic outcomes, particularly on price negotiations.

Our results are consistent with the idea that the better educational attainment and labor force participation of young women, relative to men of the same age, may lead to these women performing better in price negotiations. Women above the age of 65 are much less likely to have completed high-school or college, or to have been employed full-time when they were younger. Another subtle explanation for our results may be the effects of marriage. The data show that women above 65 are more likely to be single than those in their 30s and 40s, due to the effects of divorce as well as higher death rates for men in the same cohort. We believe it is likely that older women who buy new cars in our sample are more likely to be entering the new car market on their own for the first time, possibly having participated in it along with their spouses in the past. If so, these women may be at a disadvantage due to never having negotiated before on their own for such a big-ticket item.

References

Ayres, I. and P. Siegelman (1995, June). Race and gender discrimination in bargaining for a new car. *American Economic Review 85*(3), 304–21.

Babcock, L. R. and S. Laschever (2003). *Women Don't Ask: Negotiation and the Gender Divide*. Princeton University Press.

Bertrand, M., C. Goldin, and L. F. Katz (2010). Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, 228–255.

Bertrand, M. and K. F. Hallock (2001, October). The gender gap in top corporate jobs. *Industrial and Labor Relations Review 55*(1), 3–21.

Blau, F. D. and L. M. Kahn (2007). The gender pay gap have women gone as far as they can? *The Academy of Management Perspectives 21*(1), 7–23.

Busse, M. R., C. R. Knittel, and F. Zettelmeyer (2013). Are consumers myopic? evidence from new and used car purchases. *The American Economic Review 103*(1), 220–256.

Goldberg, P. K. (1996, June). Dealer price discrimination in new car purchases: Evidence from the consumer expenditure survey. *Journal of Political Economy 104*(3), 622–54.

Harless, D. W. and G. E. Hoffer (2002, March). Do women pay more for new vehicles? evidence from transaction price data. *American Economic Review 92*(1), 270–279.

Lang, K. and J.-Y. K. Lehmann (2012). Racial discrimination in the labor market: Theory and empirics. *Journal of Economic Literature 50*(4), 959–1006.

Langer, A. (2011). Demographic preferences and price discrimination in new vehicle sales. *manuscript, University of Michigan*.

List, J. A. (2004, February). The nature and extent of discrimination in the marketplace: Evidence from the field. *The Quarterly Journal of Economics 119*(1), 49–89.

Morton, F. S., J. Silva-Risso, and F. Zettelmeyer (2011). What matters in a price negotiation: Evidence from the us auto retailing industry. *Quantitative Marketing and Economics 9*(4), 365–402.

Morton, F. S., F. Zettelmeyer, and J. Silva-Risso (2001). Internet car retailing. *The Journal of Industrial Economics 49*(4), 501–519.

Morton, F. S., F. Zettelmeyer, and J. Silva-Risso (2003). Consumer information and discrimination: Does the internet affect the pricing of new cars to women and minorities? *Quantitative Marketing and Economics 1*(1), 65–92.

O'Neill, J. (2003). The gender gap in wages, circa 2000. *The American Economic Review 93*(2), 309–314.

Stuhlmacher, A. F. and A. E. Walters (1999). Gender differences in negotiation outcome: A meta-analysis. *Personnel Psychology 52*(3), 653–677.

Yinger, J. (1998, Spring). Evidence on discrimination in consumer markets. *Journal of Economic Perspectives 12*(2), 23–40.

Zhu, R., X. Chen, and S. Dasgupta (2008). Can trade-ins hurt you? exploring the effect of a trade-in on consumers' willingness to pay for a new product. *Journal of Marketing Research 45*(2), 159–170.