

## WHO LOSES WHEN PRICES ARE NEGOTIATED? AN ANALYSIS OF THE NEW CAR MARKET\*

AMBARISH CHANDRA<sup>†</sup>

SUMEET GULATI<sup>‡</sup>

JAMES M. SALLEE<sup>§</sup>

We establish that there are large and persistent differences in final transaction prices for identical new cars, and that demographic characteristics explain at least 20% of the observed variation. Older consumers perform progressively worse in negotiations, and the age premium is greater for women than for men. Our results suggest that the complex nature of vehicle transactions leads to price dispersion in this market, and that the worst performing groups—older women—have the lowest rates of market participation. We conjecture that the results are driven by the sharp increases in women's education and labor force participation in recent decades.

### I. INTRODUCTION

MOST MARKETS IN NORTH AMERICA HAVE FIXED, rather than negotiated, prices. Yet, two of the biggest purchases in most consumers' lives—housing and automobiles—involve negotiated prices. Both these markets are complex, involve multiple sub-negotiations, and require consumers to have considerable sophistication and experience to negotiate favorable outcomes. The large sums of money involved in these markets, combined with variation in consumers' information and bargaining abilities, likely lead to significant differences in the final prices paid.

Differences in negotiated prices should interest economists for numerous reasons. Primarily, the process of negotiating generates transaction costs

\*We are grateful for helpful suggestions from Kathy Baylis, Jim Brander, Don Fullerton, Keith Head, John Jones, Nisha Malhotra, Carol McAusland, Torrey Shanks, the Editor and the referees.

<sup>†</sup>Authors' affiliations: Department of Management, UTSC and Rotman School of Management, University of Toronto, 105 St. George Street, Toronto, Ontario Canada.  
*e-mail: ambarishchandra@gmail.com*

<sup>‡</sup>Food and Resource Economics, University of British Columbia, MCML 341, Vancouver, British Columbia, Canada  
*e-mail: sumeet.gulati@ubc.ca*

<sup>§</sup>Department of Agricultural and Resource Economics, University of California, Berkeley, 207 Giannini Hall, Berkeley, California, U.S.A.  
*e-mail: sallee@berkeley.edu*

and may lead to inefficiencies.<sup>1</sup> Second, consumers concerned about overpaying relative to a perceived ‘fair’ price may delay participating in such markets, or avoid them altogether.<sup>2</sup> Finally, negotiated markets may yield worse outcomes for consumers who are already disadvantaged or vulnerable, and may therefore exacerbate economic disparities.

Therefore, it is important to understand whether, and to what extent, there are systematic differences in negotiated outcomes. This is relevant not just for the markets highlighted above, but also for a host of other scenarios; for example, wage negotiations or bargaining terms with financial institutions. In this paper we study negotiated prices in the new car market and ask two questions: first, what is the extent of variation in final transaction prices? Second, do these prices vary systematically across demographic groups?

The new vehicle market in the U.S. offers an ideal setting for this analysis for a number of reasons. First, final prices in this market are almost always negotiated. Second, this market involves the sale of many identical goods: knowing the make, model and trim of a vehicle pins down almost exactly the good transacted. Other differences in transactions, such as the location of car dealers and the timing of the transaction, can be controlled for. Finally, we have data, from a large set of transactions, on the final prices paid by consumers, the dealer’s opportunity cost for the particular vehicle sold, and consumer demographics. Using these data we construct precise measures of the dealer’s margin in each transaction and examine how these vary according to demographics.

Following Busse and Silva-Risso [2010], we estimate equations for dealer margins controlling for a wide range of demand and supply covariates. Our data are from a major marketing firm, and contain over ten million new car transactions in the United States and Canada. Taking advantage of the large size of our sample, we study the variation in dealer margins within model, year, state and trim combinations, and across several gender and age categories. Our estimating strategy minimizes the potential impact of other factors: whether the vehicle was leased or financed; whether it was purchased at the end of the month 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

<sup>1</sup> A large theoretical literature acknowledges the importance of such transaction costs and seeks to understand why purchases such as housing and automobiles involve negotiated prices; see Bester [1993] and Wang [1995]. Recent empirical research shows that the transaction costs of negotiations in other markets can be significant; Allen *et al.* [2014] study mortgage negotiations and Jindal and Newberry [2015] examine home appliances.

<sup>2</sup> Surveys routinely show that the majority of car buyers dislike the negotiating process; a 2011 Kelly Blue Book survey showed that 59% of consumers ‘hate’ haggling, and a 2008 survey in Marketing Magazine estimates this fraction at over 80%.

stayed on dealers' lots—and whether there was a trade-in vehicle associated with the transaction.

Our results reveal large disparities in transaction prices, along with a consistent pattern of certain demographic groups paying higher prices compared to others.<sup>3</sup> The difference between the 75th and 25th percentiles of the dealer margin distribution, within a given model-trim-state combination, is almost \$1,200, even after controlling for all observable aspects of the transaction. By comparison, the average dealer margin in the data is also around \$1,200, implying that price dispersion is significant; we also show that price dispersion in this industry is considerably more than comparable measures in other industries.

We then show that at least 20% of these price differences are explained by average differences in the age and gender of consumers: older consumers pay a clear premium for new cars, and older women in particular obtain the worst outcomes. Further, the fraction of price differences explained by demographics is large compared to other variables such as the competitive environment faced by dealers or their stronger incentives to sell vehicles at month and year-end. These results persist across a range of robustness checks, and are apparent on a state-by-state basis across the U.S., as well as in the Canadian market. Revealingly, we then show that participation in the new car market is the lowest for older women, suggesting the possibility that some of these consumers avoid the new car market altogether due to their poor outcomes in negotiations.

What explains our findings? An extensive literature suggests that women fare worse in negotiations, especially wage negotiations, either due to discrimination or their own reluctance to negotiate.<sup>4</sup> However, systematic discrimination against women is an unlikely explanation for our results, given that we find younger women perform no worse than men of the same age. Differences in search and negotiation costs across various demographic groups may help explain our results. Morton *et al.* [2011] show that search costs and incomplete information have an important effect on negotiations, and that car buyers who are aware of dealers' reservation prices can capture a significant share of the dealer's margin. These factors are likely to be more important with the rise of the Internet. Savvy shoppers can infer the dealer's cost from visiting websites such as *Edmunds.com*, while consumers who do not use the Internet will be at a disadvantage in such negotiations. However, our analysis of Internet use data shows that the gender difference in Internet use does not increase with age, implying that the effects of the Internet do not constitute a complete explanation for our results.

<sup>3</sup> Note our use of the term 'disparity' does not mean that we rule out the possibility of welfare enhancing price discrimination.

<sup>4</sup> See, for example, Stuhlmacher and Walters [1999], List [2004] and Leibbrandt and List [2012].

The most likely explanation for our results is the existence of a ‘cohort effect.’ Specifically, we believe that younger female consumers negotiate better than older women due to their superior educational attainment and labor market outcomes, relative to men of the same age. The last several decades have seen dramatic improvements in socio-economic outcomes for women. Indeed, women in their twenties and thirties today have better educational outcomes on average than men, and have also succeeded in narrowing the employment and wage gap. By contrast, women in their sixties or older were far less likely to have participated in the labor force or earned a college degree when they were the same age. These differences can lead to lower information in the new car market for older women; this is especially important given the complexity of the transaction and the various margins that affect final prices including financing, monthly payments, and the trade-in allowance. Demographic data in both the United States and Canada suggest that these trends are correlated with our findings. If this hypothesis is correct, it implies that today’s cohort of young women are unlikely to do worse than their male counterparts as they age—the gender gap in negotiations may have closed, or even reversed.

In interpreting our results, we note that poor outcomes for a demographic group could be caused either by a weak bargaining position, or by weak bargaining from a given position. In standard Nash bargaining models, the outside option determines how surplus is divided. In our context, the outside option for a consumer is to buy the same model from a competing dealer, or to switch to a different car model. Thus, where we observe a demographic group paying higher than average prices for a particular model, this could be caused firstly by negotiation skills, narrowly defined: e.g., discomfort with haggling, or misunderstanding of financing details. Alternatively, the demographic group might have a worse than average outside option, due to negotiation skills, broadly defined: e.g., better awareness of alternative models, or savvy to acquire quotes from competing dealerships. Either possibility is consistent with our results. Moreover, a distinct explanation is that a given demographic group has a weak outside option, not due to skill, but rather to greater switching costs: e.g., greater brand loyalty. We cannot distinguish this from skill in our data, but note that, by comparing outcomes within the same model, we directly control for average competition in each market segment. Moreover, our geographic controls account for the availability of alternative models in each market area.

This paper is related to a number of earlier studies on the automobile industry. 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, and also 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], and failed to find evidence of worse outcomes

for women. 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.<sup>5</sup>

Thus, while there is a long tradition of research into the effect of gender in car negotiations, the role of consumers' age has been noticeably ignored.<sup>6</sup> In fact, no previous study of negotiated prices has examined how the age and gender of consumers interact.<sup>7</sup> We demonstrate that consideration of gender and age together can provide a significantly richer set of results regarding how customer demographics affect negotiations. Research focusing solely on the binary difference of gender would necessarily miss the 'cohort' effect that we suspect to be the primary driver of our results.

We further demonstrate that the interaction of age and gender can reconcile the disparate findings in the literature. Both the initial conclusion by Ayres and Siegelman [1995] that women do worse in price negotiations, and the later finding by Goldberg [1996] and Harless and Hoffer [2002] that they do not, can be generated in our data for different age cohorts. In particular, focusing on buyers aged 50 and above would suggest a clear price premium for women, but focusing on those aged under 35 would reveal virtually no gender differences.

We then show that differences in data and specification can also help explain the contrary conclusions in the prior literature. Specifically, Goldberg [1996] concludes that women do no worse than men, but we find that this may be due to the lack of detailed information on vehicle trims and options in Goldberg's data; when we ignore this information in our own study, women appear to perform significantly better, but this is due to the fact that women systematically purchase cheaper trims and options within a given car model. Similarly, when we recreate the regression specification in Harless and Hoffer [2002] by ignoring the interaction of age and gender, we find a similar result to theirs: there is no gender difference *on average*. However, this obscures the clear difference

<sup>5</sup> Other studies that examine negotiation in the automobile industry include Morton *et al.* [2001], Chen *et al.* [2008] and Langer [2016].

<sup>6</sup> Harless and Hoffer [2002] do control for customer age in the new car market, but their focus is on gender alone, and they do not disaggregate the relative performance of each gender across age groups. Langer [2016] examines the interaction between the gender and marital status of car buyers, but does not study age effects. D'Haultfoeuille *et al.* [2014] study the new car market in France, using data on consumers' age and expected income, but they lack information on gender.

<sup>7</sup> Harding *et al.* [2003] use the age of customers as a control in their study of housing transactions, but focus on gender in their analysis and do not study the interaction of age and gender.

between men and women in older age cohorts for which we find strong evidence in our full results.

Our paper also makes other contributions to the existing literature. This is the first study to document such large variation in transaction prices for identical cars, and to establish these differences with a high degree of confidence, based on a very large sample and strongly consistent results across different cuts of the data. In addition, we have detailed data on the characteristics of each vehicle sold, as well as on final transaction prices, dealers' invoice prices and transfers between manufacturers and dealers. This allows us to calculate dealer margins, rather than relying on the transaction price, which was the variable of interest in most prior work, but which includes many unobserved components. Finally, we exploit the size and high level of detail in our data to control for fine combinations of vehicle models, trims and markets, thereby alleviating concerns of unobserved interactions among these which may have affected prior studies of the role of demographics in new car sales.

Our results have important implications for public policy. The very fact that the new car market features negotiated prices implies that dealers have some ability to price discriminate. Our documented finding that final prices vary by many hundreds of dollars shows that price discrimination is in fact widespread, and that certain types of customers cross-subsidize others. While price discrimination can be efficient in some situations, it is concerning that older consumers, particularly women, pay the highest prices for new cars, given that many such consumers are likely to be retirees on fixed-incomes. It is also of concern that older women are sharply under-represented in the new car market, likely due to their worse outcomes in negotiations.

One way to reduce disparities in negotiated outcomes is to address the source of dealers' market power. Increased information can play an important role in achieving this goal. This is especially relevant for the new car and housing markets, both of which are complex transactions where consumers differ in their level of sophistication. Indeed, our results suggest that more straightforward vehicle transactions, such as those without negotiations over trade-in vehicles, have lower price dispersion and relatively better outcomes for older consumers and women. Thus, the complex nature of vehicle negotiations appears to allow dealers, who conduct many such transactions, to price discriminate at the expense of certain groups of consumers.

In Section II we present the data used in our study. In Section III we present the empirical framework. Section IV contains our main results along with various checks for the robustness of our findings. Section V discusses various explanations for our results. We conclude briefly in Section VI.

## II. DATA

We use data provided by automobile dealers to a major market research firm. These data include more than 250 key observations for each vehicle transaction including: a) vehicle characteristics: vehicle trim, number of doors, exterior color, engine type, transmission, dealer's invoice price including both factory and dealer installed accessories, and suggested retail price; b) transaction characteristics: date of the transaction, transacted price, rebates offered, how long the vehicle was on the dealer's lot, and whether the vehicle was financed, leased or a cash purchase, how much was financed, other details on financing/lease; c) trade-in characteristics: the price paid to the customer for the trade-in vehicle, and its under or over-valuation to the dealer, though we have no information on the make or model of the traded-in vehicle; d) customer demographics: gender, age, city and state of purchase. Dealers are not selected randomly, as they must agree to be included in the database. However, this is the most comprehensive dataset on new vehicle purchases that we are aware of. Our sample of the data is approximately 20% of new vehicle sales in the U.S. from April 1st, 2002 to October 31st, 2006, a total of 9,501,106 transactions (a distribution of observations across states is in the online Appendix).<sup>8</sup> As a part of our robustness exercises, we also use data, from the same source, on approximately one million transactions occurring in Canada, from May 1st, 2004 to April 15th, 2009.

Our main variables of interest are: the age and gender of the 'primary' customer, and the dealer's margin (profit) from a transaction. In joint purchases the *primary* customer is the one listed first on the invoice. Age is determined from the customer's date of birth, and gender is determined by a computer algorithm analyzing the first name.<sup>9</sup>

We calculate dealer margin as the difference between the vehicle price and the vehicle cost (also known as a dealer's invoice price).<sup>10</sup> The dealer margin from the sale of a new vehicle is determined in a complex negotiation between customer and dealer, one that has several sub-negotiations. The parties negotiate: a price of the new vehicle, a price for the vehicle

<sup>8</sup> We dropped the highest and lowest 1% of observations, according to dealer margin, from the original data obtained from the data provider. We did this to avoid using data with implausibly high or low dealer margins, but the results are similar if we use the original dataset.

<sup>9</sup> Mis-identification by the algorithm is possible, creating some measurement error in this variable.

<sup>10</sup> Vehicle price for each transaction is the price listed on contract before taxes/title fees/insurance. For our dataset this is the price that the customer pays for the vehicle and for factory and dealer installed accessories and options contracted for at the time of sale. This price is adjusted by a profit or loss made in the associated trade-in, manufacturer to dealer rebates, but not for any customer cash rebate. This price is not the Manufacturer's Suggested Retail Price (MSRP). The vehicle cost is the dealer's cost for the vehicle. For our dataset this is the retailer's 'net' or 'dead' cost for the vehicle and for both factory and dealer installed accessories contracted for at the time of sale. It includes transportation costs.

traded in, details around finance and lease conditions, the price of an extended warranty, and the choice of options (features in addition to those offered in the base). To obtain a good deal, one needs not just to ‘haggle’ well but also to understand the complex details for each sub-negotiation. Our measure of dealer margin also includes profits made (or lost) from the trade-in of the vehicle, and includes ‘holdback,’ a transfer from the manufacturer to the dealer. It does not include profits made from the sale of extended warranties.<sup>11</sup>

In this paper we focus on the dealer margin, rather than the vehicle’s price. The choice of options implies that vehicles differ in their attributes even within the same model-year and trim. The value of options is represented in both the vehicle price and vehicle cost and the difference between the two captures the outcome of this complex negotiation better than the price can.<sup>12</sup>

### II(i). *Summarizing the Data*

In Table I we present selected summary statistics across vehicle segments, using the same segment definitions as the data provider. The bottom row summarizes the entire dataset. The average vehicle sold for approximately \$27,071, generated \$1,215 in dealer margin and was on the lot for 59 days. The average customer was 46 years old, and the proportion of male buyers was 63%. Correspondingly, the median vehicle sold for \$25,384, generated \$1,010 in dealer margin, was on the lot for 28 days, and was bought by a 45 year old male customer. Luxury vehicles have the highest prices and dealer margins. Luxury, compact and sports cars stayed on dealers’ lots for fewer days than other segments. Pickup trucks have the highest proportion of male customers, while Compact cars have the highest proportion of female customers.

Measures of central tendency for our variables vary substantially even across the largest states. While females make up 34.8% of all customers in California, they make up 42% in Massachusetts. The median dealer margin is \$930 in Ohio, and \$1,135 in California.<sup>13</sup> There is also considerable variation across manufacturers. The highest selling manufacturer in our dataset was General Motors, followed by Ford, Toyota and Honda. Sports vehicle brand Porsche had the highest median vehicle price at \$56,118, the

<sup>11</sup> In two subsequent sections, the first on Robustness (IV(iii)), and then in the online Appendix we provide evidence that including profits from extended warranties does not change our results.

<sup>12</sup> The term ‘dealer margin’ should not be interpreted as a direct measure of profit. This constructed margin is occasionally negative, which does not necessarily mean that dealers lose money on the transaction, although that is possible. Instead, there are other reasons why dealers may record negative margins, including responding to manufacturer provided incentives; clearing out cars with low demand; taking losses on the new car in order to make profits on the trade-in; or generating future servicing incomes.

<sup>13</sup> We illustrate these differences in more detail in the online Appendix.



TABLE I  
STATISTICS ACROSS SEGMENTS

Segment	Mean						Median					
	Male Cust	Female Cust	Veh Price	Dealer Marg	Turn Days	Cust Age	Veh Price	Dealer Marg	Turn Days	Cust Age		
Compact (16%)	52.6%	47.4%	\$17,133	\$852	53.0	44.3	\$16,463	\$705	25	44		
Large (<1%)	72.8%	27.2%	\$26,413	\$1,080	84.5	66.6	\$26,390	\$977	53	68		
Luxury (8%)	64.2%	35.8%	\$39,411	\$1,694	46.4	49.1	\$35,531	\$1,512	20	49		
Midsized (18%)	55.7%	44.3%	\$23,023	\$1,011	57.6	48.6	\$22,284	\$865	29	48		
Pickup (17%)	81.7%	18.3%	\$28,343	\$1,269	69.0	44.6	\$27,802	\$1,106	39	44		
SUV (28%)	61.3%	38.7%	\$31,010	\$1,342	58.8	44.8	\$29,418	\$1,152	28	44		
Sporty (4%)	64.0%	36.0%	\$27,529	\$1,485	61.6	42.7	\$24,922	\$1,186	26	43		
Van (6%)	68.6%	31.4%	\$26,995	\$1,245	61.6	46.6	\$26,657	\$1,036	28	43		
Total (100%)	63.3%	36.7%	\$27,071	\$1,215	58.8	45.9	\$25,384	\$1,010	28	45		

Notes: Source: Authors' Calculations

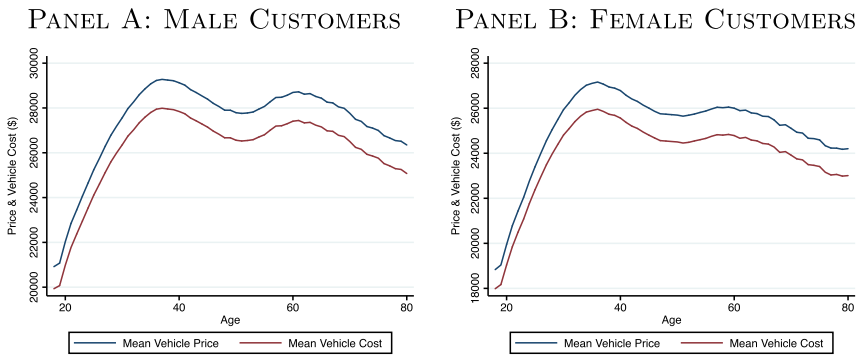


Figure 1  
 Vehicle Price and Cost by Age (21-80 years) [Colour figure can be viewed at [wileyonlinelibrary.com](http://wileyonlinelibrary.com)]

highest dealer margin at \$3,228, and the highest proportion of male customers at 77.8%. Hyundai sold vehicles with the lowest median vehicle price at \$18,402, the lowest dealer margin at \$532, and the highest proportion of female customers at 47.1%.

The average transaction price, vehicle cost, and dealer margin also vary substantially by age and gender. We illustrate this in Figure 1 where we plot the mean transacted price and vehicle cost separately for each gender, among customers under 80 years old. For both genders the average transaction prices, as well as the average vehicle costs, are highest among 36–37 year olds and lowest among 18-year olds. Prices and costs have a twin-peak distribution with the second peak occurring at age 61 for men and 59 for women. On average, men purchase vehicles that are about \$2,000 more expensive, according to both transaction price and vehicle cost, than those bought by women of the same age.

Dealer margins follow a similar distribution across age groups, which is not surprising as margins tend to rise with the transaction price. However, if we plot dealer margin as a percentage of vehicle cost an interesting pattern emerges (see figure 2). Among male customers, dealer margins as a percentage of vehicle cost first fall with age and then rise. Among female customers, there is only a slight initial decline, but generally average dealer margins as a percentage of vehicle cost tend to rise with the age of the customer.

When we examine the proportion of male and female customers across age groups, we observe a pattern of ‘missing women’ (see columns titled Row % in Table II). From comprising almost 50% of transactions among under 25 year-olds, women comprise only 32% of all transactions among consumers over 70. This is despite the fact that, in the overall population, women start to outnumber men after about age 30 and that, after age 70, women make up about 60% of the population. Moreover, the proportion of licensed drivers is similar at every age for men and women, as revealed

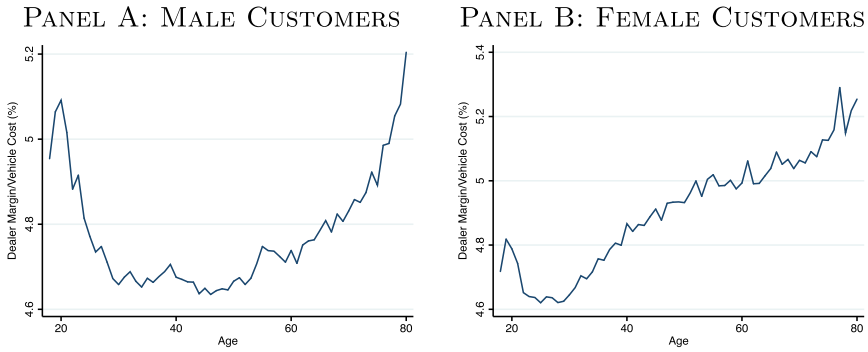


Figure 2

Dealer Margin as a Percentage of Vehicle Cost by Age (21-80 years) [Colour figure can be viewed at [wileyonlinelibrary.com](http://wileyonlinelibrary.com)]

by data from the Federal Highway Administration. This pattern of declining shares of female consumers with age is intriguing, and is likely to be related to our findings; we return to this issue in Section V.

III. EMPIRICAL STRATEGY

Our goal is to examine whether there are systematic differences in the performance of various demographic groups when negotiating prices in the new car market. We can do this in a comprehensive manner given the high level of detail in the data: each transaction records the make, model, model-year and trim of the vehicle, as well as the date of the sale and other features of the transaction.

TABLE II  
SALES BY AGE CATEGORY

Age Category	Gender								
	Male			Female			Total		
	No.	Col %	Row %	No.	Col %	Row %	No.	Col %	Row %
Age Under 25	279,721	4.7	50.7	271,921	7.8	49.3	551,642	5.8	100.0
Age 25-30	440,851	7.3	57.6	324,133	9.3	42.4	764,984	8.1	100.0
Age 30-35	606,300	10.1	63.1	354,206	10.2	36.9	960,506	10.1	100.0
Age 35-40	703,949	11.7	64.7	384,088	11.0	35.3	1,088,037	11.5	100.0
Age 40-45	781,372	13.0	63.8	444,170	12.7	36.2	1,225,542	12.9	100.0
Age 45-50	785,989	13.1	63.1	458,736	13.2	36.9	1,244,725	13.1	100.0
Age 50-55	709,303	11.8	63.5	406,902	11.7	36.5	1,116,205	11.8	100.0
Age 55-60	599,316	10.0	65.7	313,537	9.0	34.3	912,853	9.6	100.0
Age 60-65	410,178	6.8	67.0	201,835	5.8	33.0	612,013	6.4	100.0
Age 65-70	276,398	4.6	68.1	129,380	3.7	31.9	405,778	4.3	100.0
Age Over 70	414,755	6.9	67.6	199,164	5.7	32.4	613,919	6.5	100.0
<b>Total</b>	<b>6,008,132</b>	<b>100.0</b>	<b>63.3</b>	<b>3,488,072</b>	<b>100.0</b>	<b>36.7</b>	<b>9,496,204</b>	<b>100.0</b>	<b>100.0</b>

Notes: Source: Authors' Calculations

We do not have direct information on certain options and after-market purchases which may affect the transaction price. However, the information on these options and accessories is embedded in our data on the vehicle's invoice price, which the dealer records and reports to our data provider. Therefore, if a customer purchases a certain after-market option, such as a ski rack or all-weather tires, the value of this option will be added to both the transaction price and the dealer's invoice price. For this reason, examining the *dealer's margin*, as we can do in our data, is superior to examining either the transaction price or the discount from the manufacturer's suggested price, which were employed in previous research (for example, Goldberg [1996] and Langer [2016]). If certain demographic groups generate systematically different dealer margins, holding constant the attributes of the purchased vehicle and other features of the transaction, then we can confidently attribute these differences to the negotiating process.

We keep in mind two goals. The first is to examine dealer margins within as uniform a product as possible. Our very large sample size enables us to examine variation within a combination of vehicle model, model-year, trim and state. This is helped by the presence of many high-selling vehicle models, which have over 100,000 transactions during our sample period (some have almost 300,000 transactions). These popular models draw customers from a wide range of gender and age groups and therefore allow us to identify the variation across gender and age combinations while keeping constant the characteristics of the vehicle. Figure 3 shows the customer age distributions, separately for each gender, for two of the most popular car models during our sample period—the Honda Accord and the Toyota Camry. Both models are purchased by substantial numbers of consumers in each age and gender category, which enables our empirical identification strategy.<sup>14</sup>

The second goal is to minimize the potential bias from omitted variables. Recent research indicates that besides model and dealer characteristics, transaction characteristics influence price and dealer profits. Whether the vehicle was leased or financed, or whether 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 gender and age categories exhibit systematic variation over transaction characteristics, omitting them would bias our results. Such variables, or suitable proxies, are included as controls.

We adapt a reduced form estimating equation from Busse and Silva-Risso [2010] to estimate a simple linear model of dealer profits from a

<sup>14</sup> Also see the online Appendix for the full distribution of purchases of each age and gender grouping across vehicle segments.



Figure 3  
Customer Age Distribution for Popular Models

particular vehicle transaction. We denote a model-trim combination by  $m$ , model-year by  $y$ , state by  $s$ , and date by  $t$ . Then Dealer  $d$ 's profits from a transaction with customer  $i$  is expressed as:

$$(1) \quad \pi_{myst}^{id} = \beta_0 + \beta_1 \text{demog}^i + \beta_2 \text{transchar}_t^{id} + \lambda_d + \mu_t + \Theta_{mys} + \epsilon_t^{id}$$

Specifically, the dealer's profit depends on a) the demographics of the customer, captured by age and gender; b) characteristics of the transaction—which includes timing-related measures such as whether the vehicle sold on a weekend, the end of a month or the end of a year—and characteristics of the purchase such as whether the vehicle was leased, financed or bought with cash; c) dealer fixed effects, denoted  $\lambda_d$ , which will be employed in a robustness exercise; d) year-month fixed effects, denoted  $\mu_t$ ; e) characteristics of the vehicle, denoted  $\Theta_{mys}$ . Estimating this equation yields the average difference in dealer margin within model-state-year-trim across our gender and age categories.

We emphasize our ability to exploit the data to control flexibly for  $\Theta_{mys}$ . It is not enough to control simply for the model and model-year being purchased by the consumer, as is done in a number of earlier studies, for example Goldberg [1996] and Langer [2016]. This is because, even within a given model, customers can choose different trims in a manner that may be

correlated with their demographics.<sup>15</sup> But even controlling for the model and trim, as has been done in other studies, for example Harless and Hoffer [2002], may not be enough. This is partly because there may be other options that consumers purchase, which we capture by examining dealer margins, but also because prices or margins for a given model-trim combination are likely to vary across markets due to differences in consumer demand or the network of dealers.

To address all of these possible issues, we control in our regressions for the combination of a model, model-year, trim and the state in which it was purchased, and estimate demographic differences within these combinations. As a concrete example, we compare consumers purchasing a 2007 Honda Civic LX in California with other consumers purchasing exactly the same vehicle in the same state, while also accounting for unobserved options in each purchase. In additional robustness checks we include fixed effects for the city and state in which the purchaser resides. No prior study of demographics in new car sales has accounted for such a fine level of variation across transactions. We argue that, once all of these features of various transactions are accounted for, differences in dealer margins must emerge from the negotiation process rather than from different vehicle choices by consumers, as we discuss in more detail below.

#### IV. RESULTS

We now present our results. This section is divided into four parts. We first document the extent of variation in new car prices, showing that there are very large differences in final prices paid for the same new car, even after controlling for all observable characteristics of the transaction. Next, we show that the demographics of consumers explain a significant portion of these differences. We then show that these results do not change in response to a wide variety of robustness checks, and that our results also extend to the Canadian market. Finally, we establish the effects of competition on prices and price dispersion, and relate these to our results on demographics.

##### IV(i). *Variation in Dealer Margins*

In this subsection we establish that there is huge variation in negotiated prices for new vehicles. Differences in the final prices for the same new car model can be many hundreds of dollars, even accounting for all observable aspects of the transaction.

We emphasize again that our main variable of interest is the dealer's margin on the new car, which is recorded as the difference between the

<sup>15</sup> Indeed, we will show later that this is precisely the case, as women tend to buy not just cheaper models than men, but also cheaper trims within a given model.

TABLE III  
VARIATION IN DEALER MARGINS

Sample	Dealer Margin		Residuals-1		Residuals-2	
	90 to 10	75 to 25	90 to 10	75 to 25	90 to 10	75 to 25
Full Sample	3,048	1,463	2,570	1,214	2,515	1,187
No Trade-ins	2,989	1,423	2,458	1,158	2,398	1,129
No Financing	3,167	1,543	2,565	1,220	2,501	1,184
No Leases	2,967	1,420	2,515	1,184	2,461	1,160
California	3,366	1,645	2,889	1,364	2,816	1,331
Texas	3,005	1,443	2,620	1,249	2,580	1,229

*Notes:* 90 to 10 refers to the difference between the 90th and 10th percentile in each distribution; analogously for 75 to 25. Residuals-1 refers to the distribution of residuals from a regression of dealer margins on model\*year\*trim\*state fixed effects. Residuals-2 adds year-month fixed effects, city-state fixed effects and other controls to the regression used to generate residuals.

purchase price and the dealer's invoice price, for reasons discussed in Section III. In Table III, we summarize two measures of dispersion in dealer margins. We focus on the difference between the 90th and 10th percentiles of the dealer margin distribution, as well as between the 75th and 25th percentiles. Examining these percentiles allows us to consistently measure and compare the dispersion in prices paid across various subsamples of the data, while also ignoring outlier observations that can distort such comparisons.

The first line of the Table shows these measures of dispersion for the full sample of almost 10 million observations. The difference in dealer margins between the 90th and 10th percentiles is over \$3,000, while the interquartile range is over \$1,400.<sup>16</sup> These differences may appear large, but it can be misleading to make comparisons over the entire distribution of prices ignoring other features of the transaction. Therefore, in successive columns of Table III we control for observable characteristics. In columns 3 and 4 we examine the dispersion in the distribution of residuals from a regression of dealer margins on fixed effects for each combination of model-trim, model-year and state. In columns 5 and 6 we add to this regression a number of other controls: fixed-effects for the year-month of the transaction and for the city-state where the buyer resides; controls for whether the vehicle was financed, leased, or was a cash payment; and indicators for whether the transaction took place on a weekend, or at the end of the month or year, when dealers have incentives for higher sales.

The dispersion among the residuals is naturally smaller than in the full distribution of dealer margins. Nevertheless, even after controlling for all observable aspects of the transaction, the difference between the 90th and 10th percentiles of the dealer margin distribution is about \$2,500, and the interquartile range is almost \$1,200. In comparison to

<sup>16</sup> By comparison, the average dealer margin in our data is approximately \$1,200.

TABLE IV  
COMPARISON OF PRICE DISPERSION MEASURES

Study: Distribution:	Chandra & Tappata [2011]	Allen et al [2014]		This paper	
	Margins	Margins	Residuals	Margins	Residuals
C.V.	0.25	0.47	0.39	1.01	0.86
IQR/Mean	0.35	0.66	0.52	1.20	0.97

Notes: Margins defined as the difference between the final price and the proxy for marginal cost. Residuals constructed from a regression of margins on a range of covariates. C.V. is the ratio of the Standard Deviation to the Mean.

the values reported in Table I, this implies that the interquartile range difference is, on average, 100% of the dealer margin and about 5% of the transaction price. In subsequent rows of Table III we show that these large differences in dealer margins persist across various subsamples of the data. In particular, we drop transactions in which customers traded in older vehicles, as well as those involving financing from the dealer or leased vehicles. We also restrict the sample to the two largest states—California and Texas. In all cases, the dispersion in dealer margins remains very high; always well over a thousand dollars for the interquartile range even after controlling for attributes of the vehicle and other observable features of the transaction.

The extent of price dispersion that we observe in automobile transactions is large compared to other settings. To show this, we compare our measures of price dispersion to those in two other markets: mortgage interest rates—which also involve negotiated prices—and retail gasoline, which has fixed prices but has been shown by prior research to exhibit high spatial and temporal price dispersion. These industries are also attractive for this comparison because of the availability of data on firm margins, which are generally difficult to measure accurately. Both of these industries have reliable data on the most important proxy for firms' marginal costs: the spot price of gasoline in the case of gas stations, and the five-year bond rate in the case of mortgage rates.

The mortgage market is, in fact, an excellent setting in which to study negotiated prices, since most in-branch loan officers have considerable discretion to lower rates from those posted by national lenders. As a result, new homebuyers often solicit multiple offers and bargain extensively with banks. For our purposes we summarize the measures of price dispersion in this industry reported in Allen *et al.* [2014]. For retail gasoline, we summarize measures of price dispersion using the data in Chandra and Tappata [2011], which we obtained for this purpose. Table IV compares these three studies according to two commonly used measures of price dispersion: the coefficient of variation, which is defined as the ratio of the standard deviation to the mean, and the ratio of the inter-quartile range to the mean.



Both measures are unit-less and therefore directly comparable across the three studies. In addition, the Allen *et al.* [2014] study used a similar method to ours of constructing the residuals of firm margins by controlling for a rich set of covariates; we therefore present price dispersion measures for this distribution of residuals as well.

As Table IV shows, price dispersion in our data on automobile transactions is considerably higher than in the other two industries, according to either measure of dispersion. The dispersion in gasoline, while high in comparison to many other fixed-price markets, is easily the lowest among the three industries, which highlights the role of negotiated prices in creating price dispersion. But even in the comparison between the two negotiated price markets, automobile margins are around twice as dispersed as margins in the home mortgage industry.

#### IV(ii). *Regression Results*

We now turn to our regression results. Before estimating our main specification in Equation (1), we first show how various demographic groups choose new vehicles of differing value. To do this, we examine the dealer's invoice price of the vehicle, which we hereafter refer to as the *vehicle cost*. The vehicle cost is determined in advance of the transaction and so does not reflect any elements of the negotiating process. Therefore, differences in vehicle costs across demographic groups must purely reflect different choices by consumers.

Table V presents the results from regressing the log of the vehicle cost in each transaction on demographic characteristics and other controls. The omitted demographic group is male consumers under the age of 25. In column 1 we control simply for the state in which the transaction took place, as well as year\*month fixed effects. The results show that the cheapest vehicles are bought by women under 25, who pay about 8% less than men of the same age. The most expensive vehicles are purchased by 35–40 year old men who pay 20% more than the omitted category.

Differences in vehicle costs can be driven either by the consumer's choice of model, or by the choice of more expensive trims within a model. Therefore, in column 2 we control for the combination of state, model and model-year. This naturally produces much smaller differences, since any variation now must be driven almost entirely by the choice of vehicle trim. Nevertheless, we see a significant relationship between gender and the choice of trim. Within a given model, women appear to consistently buy cheaper trims than men of the same age; the difference is, on average, about 0.7% of the cost of the vehicle, which is statistically significant in almost all age groups, and which translates to about \$200 on average. In column 3 we estimate fixed-effects for the interaction of state, model, model-year and trim. Differences between the genders now are tiny and generally not

TABLE V  
REGRESSION OF LOG(VEHICLE COST)

	(1)	(2)	(3)
Age < 25 Female	-0.084 <sup>a</sup> (0.006)	-0.003 <sup>a</sup> (0.001)	0.004 <sup>a</sup> (0.001)
Age 25-30 Male	0.109 <sup>a</sup> (0.005)	0.009 <sup>a</sup> (0.001)	0.009 <sup>a</sup> (0.001)
Age 25-30 Female	0.039 <sup>a</sup> (0.006)	0.002 <sup>b</sup> (0.001)	0.007 <sup>a</sup> (0.001)
Age 30-35 Male	0.180 <sup>a</sup> (0.008)	0.015 <sup>a</sup> (0.001)	0.012 <sup>a</sup> (0.001)
Age 30-35 Female	0.114 <sup>a</sup> (0.008)	0.007 <sup>a</sup> (0.001)	0.010 <sup>a</sup> (0.001)
Age 35-40 Male	0.209 <sup>a</sup> (0.009)	0.017 <sup>a</sup> (0.001)	0.014 <sup>a</sup> (0.001)
Age 35-40 Female	0.131 <sup>a</sup> (0.008)	0.010 <sup>a</sup> (0.001)	0.011 <sup>a</sup> (0.001)
Age 40-45 Male	0.194 <sup>a</sup> (0.008)	0.015 <sup>a</sup> (0.001)	0.012 <sup>a</sup> (0.001)
Age 40-45 Female	0.108 <sup>a</sup> (0.007)	0.009 <sup>a</sup> (0.001)	0.010 <sup>a</sup> (0.001)
Age 45-50 Male	0.165 <sup>a</sup> (0.007)	0.011 <sup>a</sup> (0.001)	0.010 <sup>a</sup> (0.001)
Age 45-50 Female	0.084 <sup>a</sup> (0.007)	0.008 <sup>a</sup> (0.001)	0.010 <sup>a</sup> (0.001)
Age 50-55 Male	0.155 <sup>a</sup> (0.007)	0.011 <sup>a</sup> (0.001)	0.009 <sup>a</sup> (0.001)
Age 50-55 Female	0.079 <sup>a</sup> (0.007)	0.008 <sup>a</sup> (0.001)	0.010 <sup>a</sup> (0.001)
Age 55-60 Male	0.171 <sup>a</sup> (0.008)	0.013 <sup>a</sup> (0.001)	0.010 <sup>a</sup> (0.001)
Age 55-60 Female	0.085 <sup>a</sup> (0.008)	0.009 <sup>a</sup> (0.001)	0.010 <sup>a</sup> (0.001)
Age 60-65 Male	0.180 <sup>a</sup> (0.009)	0.013 <sup>a</sup> (0.001)	0.010 <sup>a</sup> (0.001)
Age 60-65 Female	0.080 <sup>a</sup> (0.009)	0.006 <sup>a</sup> (0.001)	0.009 <sup>a</sup> (0.001)
Age 65-70 Male	0.165 <sup>a</sup> (0.010)	0.009 <sup>a</sup> (0.001)	0.007 <sup>a</sup> (0.001)
Age 65-70 Female	0.064 <sup>a</sup> (0.009)	0.001 (0.001)	0.006 <sup>a</sup> (0.001)
Age > 70 Male	0.114 <sup>a</sup> (0.011)	-0.003 <sup>b</sup> (0.001)	0.002 <sup>c</sup> (0.001)
Age > 70 Female	0.020 <sup>b</sup> (0.010)	-0.006 <sup>a</sup> (0.001)	0.004 <sup>a</sup> (0.001)
Financed Indicator	-0.066 <sup>a</sup> (0.005)	-0.006 <sup>a</sup> (0.000)	-0.000 (0.000)
Leased Indicator	0.107 <sup>a</sup> (0.009)	-0.001 (0.001)	0.005 <sup>a</sup> (0.001)
Sat or Sun FE	-0.023 <sup>a</sup> (0.001)	-0.003 <sup>a</sup> (0.000)	-0.001 <sup>a</sup> (0.000)
End of Month FE	0.016 <sup>a</sup> (0.001)	0.001 <sup>a</sup> (0.000)	0.000 (0.000)
End of Year FE	-0.013 <sup>a</sup> (0.002)	-0.003 <sup>a</sup> (0.000)	-0.002 <sup>a</sup> (0.000)
Constant	10.013 <sup>a</sup> (0.025)	10.072 <sup>a</sup> (0.005)	10.070 <sup>a</sup> (0.005)
R <sup>2</sup>	0.096	0.884	0.930

Notes: <sup>c</sup> $p < 0.1$ , <sup>b</sup> $p < 0.05$ , <sup>a</sup> $p < 0.01$ . Fixed-effects specified as Col. 1: state; Col. 2: model\*model-year\*state; Col. 3: model\*model-year\*trim\*state. All regressions include year\*month FE's. Robust standard errors clustered by model\*model-year in parentheses. N=9,496,204.

significant. Any remaining differences must be due to consumers' choice of accessories and aftermarket options, which will be controlled for in our regressions of dealer margin.

These results suggest that there are clear gender differences in the choice of models, as well as trims within a model, which are important to control for in order to credibly establish differences in negotiating patterns across demographic groups, which is what we turn to now in our estimation of Equation (1).

Table VI presents results from a set of regressions where the dependent variable is the dealer's margin on each transaction and the main regressors of interest are the age and gender of consumers.<sup>17</sup> The omitted age group is male consumers below 25. Column 1 includes fixed-effects for the State of the transaction, Column 2 uses model\*model-year\*state fixed-effects,

<sup>17</sup> We do not transform the dealer margin by taking a log, as this would require dropping observations where margins are negative. Later, we show that our results are robust to using only transactions with positive dealer margins or to expressing the dependent variable as a fraction of the vehicle cost.

TABLE VI  
REGRESSION OF DEALER MARGIN

	(1)	(2)	(3)
Age < 25 Female	-137.6 <sup>a</sup> (8.3)	-38.2 <sup>a</sup> (4.2)	-23.0 <sup>a</sup> (3.3)
Age 25-30 Male	73.1 <sup>a</sup> (8.5)	-28.4 <sup>a</sup> (3.5)	-27.7 <sup>a</sup> (3.2)
Age 25-30 Female	-31.7 <sup>a</sup> (9.5)	-44.4 <sup>a</sup> (4.1)	-32.5 <sup>a</sup> (3.5)
Age 30-35 Male	143.7 <sup>a</sup> (17.0)	-43.5 <sup>a</sup> (4.4)	-44.2 <sup>a</sup> (3.8)
Age 30-35 Female	68.3 <sup>a</sup> (15.4)	-26.1 <sup>a</sup> (4.4)	-17.9 <sup>a</sup> (3.7)
Age 35-40 Male	192.4 <sup>a</sup> (19.5)	-34.0 <sup>a</sup> (4.5)	-36.5 <sup>a</sup> (4.0)
Age 35-40 Female	120.7 <sup>a</sup> (17.0)	2.6 (4.4)	8.6 <sup>b</sup> (3.9)
Age 40-45 Male	179.1 <sup>a</sup> (16.7)	-30.3 <sup>a</sup> (4.2)	-31.3 <sup>a</sup> (3.6)
Age 40-45 Female	121.8 <sup>a</sup> (14.8)	27.2 <sup>a</sup> (4.3)	34.1 <sup>a</sup> (3.7)
Age 45-50 Male	151.0 <sup>a</sup> (13.3)	-22.4 <sup>a</sup> (4.3)	-21.3 <sup>a</sup> (3.5)
Age 45-50 Female	112.9 <sup>a</sup> (12.3)	49.1 <sup>a</sup> (4.5)	56.7 <sup>a</sup> (3.7)
Age 50-55 Male	147.7 <sup>a</sup> (12.1)	-6.3 (4.4)	-5.4 (3.5)
Age 50-55 Female	119.2 <sup>a</sup> (12.4)	69.1 <sup>a</sup> (4.7)	76.1 <sup>a</sup> (3.8)
Age 55-60 Male	185.2 <sup>a</sup> (13.1)	24.9 <sup>a</sup> (4.6)	24.5 <sup>a</sup> (3.6)
Age 55-60 Female	139.0 <sup>a</sup> (13.6)	92.7 <sup>a</sup> (4.9)	99.5 <sup>a</sup> (4.0)
Age 60-65 Male	201.4 <sup>a</sup> (14.7)	41.1 <sup>a</sup> (4.5)	40.3 <sup>a</sup> (3.6)
Age 60-65 Female	144.4 <sup>a</sup> (14.9)	113.0 <sup>a</sup> (5.1)	121.9 <sup>a</sup> (4.2)
Age 65-70 Male	196.0 <sup>a</sup> (15.5)	72.7 <sup>a</sup> (4.7)	75.1 <sup>a</sup> (4.0)
Age 65-70 Female	141.7 <sup>a</sup> (15.9)	138.4 <sup>a</sup> (5.9)	150.3 <sup>a</sup> (5.0)
Age > 70 Male	178.9 <sup>a</sup> (16.2)	135.1 <sup>a</sup> (5.1)	146.8 <sup>a</sup> (4.5)
Age > 70 Female	119.0 <sup>a</sup> (16.9)	179.3 <sup>a</sup> (6.5)	199.0 <sup>a</sup> (5.5)
Financed Indicator	-36.5 <sup>b</sup> (15.0)	127.4 <sup>a</sup> (3.4)	135.8 <sup>a</sup> (3.3)
Leased Indicator	217.2 <sup>a</sup> (18.2)	181.0 <sup>a</sup> (8.4)	194.5 <sup>a</sup> (8.1)
Sat or Sun FE	-21.0 <sup>a</sup> (3.0)	14.0 <sup>a</sup> (1.4)	17.1 <sup>a</sup> (1.3)
End of Month FE	-67.4 <sup>a</sup> (2.4)	-84.8 <sup>a</sup> (1.5)	-84.8 <sup>a</sup> (1.5)
End of Year FE	-55.3 <sup>a</sup> (4.8)	-26.9 <sup>a</sup> (3.7)	-24.6 <sup>a</sup> (3.7)
Constant	996.7 <sup>a</sup> (35.7)	81.6 (50.4)	50.6 (51.4)
R <sup>2</sup>	0.022	0.245	0.264

Notes: <sup>c</sup>  $p < 0.1$ , <sup>b</sup>  $p < 0.05$ , <sup>a</sup>  $p < 0.01$ . Fixed effects specified as Col. 1: state; Col. 2: model\*model-year\*state; Col. 3: model\*model-year\*trim\*state. All regressions include year\*month FEs. Robust standard errors clustered by model\*model-year in parentheses. N=9,496,204.

and Column 3 uses model\*model-year\*trim\*state fixed-effects. All regressions also include year\*month fixed-effects to control for general trends in vehicle costs and prices. Standard errors are clustered by model\*model-year.<sup>18</sup>

The results in the first column show that dealers, on average, make \$137 less from female consumers under the age of 25 than from male consumers of the same age. Male consumers between 60 and 65 years generate approximately \$200 more in margins than the omitted group of male consumers under 25, and about \$340 more than female consumers under 25.

Note, however, that since Column 1 does not control for the make, model or trim of the vehicle, the results for different age and gender groups may be driven by the types of cars they purchase. For example, if young, female consumers tend to buy cheaper models or trims, and if dealers

<sup>18</sup> This appeared to be the most natural level at which to cluster standard errors. Clustering at the model level, instead, has no appreciable impact on the standard errors.




Figure 4  
Dealer Margin across Demographic Groups, Relative to Men under 25

make lower margins on cheaper cars, then the results may just reflect these choices rather than differences in negotiation. Therefore, in column 2 of Table VI, we estimate the same relationship as before, but we include model\*model-year\*state fixed-effects. We now find that there appears to be a clear age premium in new car negotiations—the youngest consumers pay considerably less than older consumers for identical vehicles.<sup>19</sup> We also find that, with the exception of the youngest age category of consumers under 25, men negotiate lower prices on average than women of the same age. These differences are small at first, but grow steadily among older consumers. In column 3, we estimate this relationship within a vehicle's model-trim. We find that the differences persist, and are in fact somewhat larger among older consumers.

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 illustrate this in Figure 4, by plotting the coefficient estimates from Column 3 of Table VI. The shaded regions around each plot are the 95% confidence intervals. F-tests indicate that the coefficients for men and women of the same age are

<sup>19</sup> This finding corroborates similar results in Harless and Hoffer [2002] and D'Haultfoeuille *et al.* [2014].

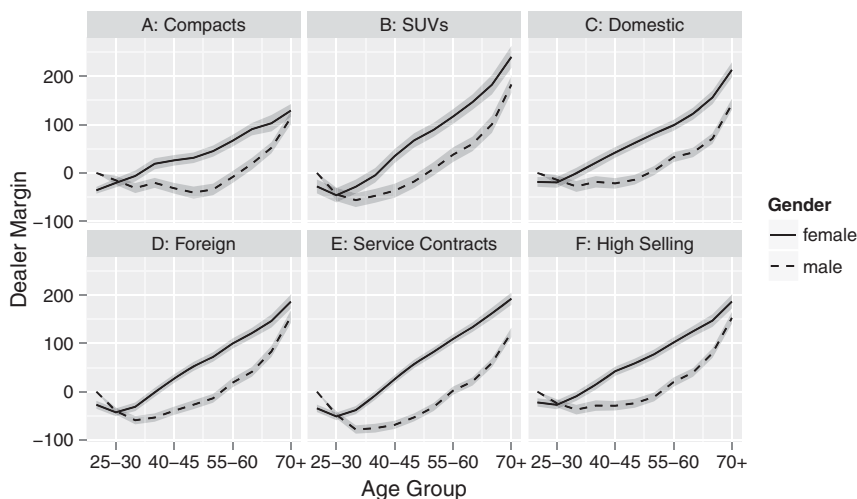


Figure 5  
Dealer Margins for Subsamples, by Gender and Age

almost always significantly different from each other at the 1% level. (The exception is the 25-30 age group.)

According to the coefficients from column 3, dealer margins are the lowest for 30-35 year old men, who pay \$36 less than the omitted category of men under 25, and highest for women over 70 who pay \$200 more. Therefore, there is a \$236 difference, on average, between the highest and lowest paying demographic groups. The average dealer margin earned across all vehicles in our data is \$1,215 (see Table I). Thus, the oldest women generate a margin that is around 20% higher than the lowest paying consumers who buy the same vehicle.<sup>20</sup> Note that these differences are driven by variation in negotiating outcomes for a given model-trim, since there is no evidence that elderly women, for example, are being manipulated into purchasing higher trim levels by car salesmen.

Turning to the other factors that may influence dealer margins, we see from the third column in Table VI that dealers make \$135-\$195 higher margins on cars that are leased or financed, relative to outright cash purchases. This accords with common observations about the new car market. Finally, dealers clearly earn lower margins at the end of the month and year. On average, their margins are lower by \$85 and \$109 (the sum of the end-of-month and end-of-year coefficients) at these times, again conforming to casual observations that dealers respond to manufacturer incentives

<sup>20</sup> This is an underestimate of the true difference across demographic groups because of our choice of five-year groupings of age categories; finer age categories would lead to larger differences between the extremes.

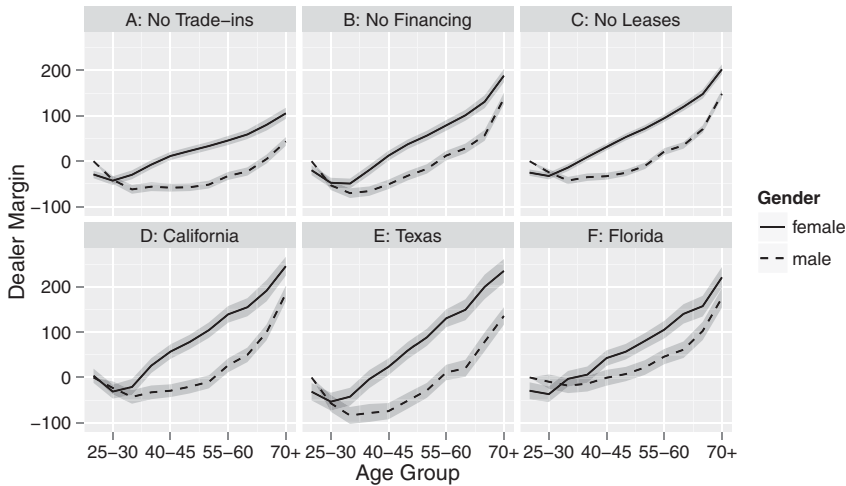


Figure 6  
Dealer Margins for Subsamples, by Gender and Age

at the end of calendar months and years. Note that these lower margins may be caused either by dealers being willing to sacrifice profits to meet sales quotas, or by more price-conscious consumers choosing to purchase cars at these times, knowing the incentives that dealers face.

#### IV(iii). *Robustness*

In this section we show that our results are not driven by omitted variables, outlier observations, or selection. We present the main coefficients of interest graphically in Figures 5 and 6. In each figure, the solid line represents dealer margins on female customers relative to men under 25, while the dashed line corresponds to margins on male customers. These results are obtained from regressions using the same specification as Column 3 of Table VI, which include fixed effects for model-year-states-trim and for year-months.<sup>21</sup>

We start by examining particular segments—Compact cars and SUV’s in Panels A and B of Figure 5. These segments are disproportionately associated with certain age groups; younger consumers are more likely to purchase small cars, while middle-aged consumers, especially those who are married, are more likely to purchase family cars such as SUV’s. We then examine domestic (North American) and foreign manufacturers separately in panels C and D. Anecdotal reports suggest that dealers of foreign cars

<sup>21</sup> Regression results, and the full set of clustered standard errors for these exercises, are presented in the online Appendix.

are less likely to negotiate on final prices; if so, this may affect our results if demographic groups differ in their propensity to purchase foreign cars.

In Panel E of Figure 5 we include dealer profits from service contracts in our definition of dealer margins. Our data indicate that these service contracts are disproportionately likely to be purchased by older buyers, and therefore may affect negotiated prices and the recorded dealer margin if dealers are willing to sacrifice profits on the new car sale, expecting to make these up on the sale of service contracts.<sup>22</sup> Note, though, that such behavior would strengthen our earlier results.

We then show, in Panel F, that our results are robust to examining popular car models alone; this is to verify that our results 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 50,000 units sold during our sample period. This limits our sample to 53 models, out of an initial set of more than 600 models, but they comprise around 56% of total 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 a higher payment on some trade-ins in order to make greater margins on the new car, or vice versa.<sup>23</sup> We therefore restrict attention, in Panel A of Figure 6, to the subset of transactions where consumers did not trade-in an older vehicle.

We then control for the possibility that financed or leased vehicles may affect the results, even though we have included controls for such transactions in our regressions. Harless and Hoffer [2002] dropped leased cars from their sample of vehicle transactions, arguing that there is generally little room to negotiate prices in such cases, although both media reports and an examination of our data suggest that this is not the case. Additionally, transactions that involve financing through the dealer may be problematic if, for example, the dealer is willing to accept a lower price on the car in return for higher interest payments. Therefore, in Panels B and C of Figure 6 we drop observations which involve financed and leased vehicles, respectively.

Panels D, E and F of Figure 6 present the results for the three largest states in the country and in our sample—California, Texas and Florida—to show that the results hold across geographic regions of the country.

In general, most of these robustness checks convey the same basic result as in our main specification with the full sample: there is a clear age premium in new car sales, and a steeper rise with age for women than for

<sup>22</sup> Sallee [2011] cites the purchasing of a service contracts as evidence that a car buyer is less savvy.

<sup>23</sup> See Zhu *et al.* [2008] and Busse and Silva-Risso [2010] for evidence regarding this possibility.

men. In some sub-samples, such as for compact cars, foreign cars, and high-selling cars, the gender difference narrows at the highest age categories, but it remains the case that a significant gender gap opens up among middle-aged buyers even if this is not always maintained among the oldest consumers.<sup>24</sup>

Among all these robustness checks the smallest age-gradient appears in Panel A of Figure 6. The panel shows that, when we exclude trade-ins, the premium for the oldest age groups is quite small for women, at about \$100, and almost non-existent for men. This is revealing; it suggests that transactions in which consumers do not trade-in old vehicles leave less room for dealers to extract profits on the sale of the new car. We believe that this piece of evidence fits well with the larger explanation that the complexity of vehicle transactions fosters price dispersion that benefits some consumers at the expense of others. We will return to this issue in Section V.

In additional robustness checks we test for a number of further possibilities that may affect the results. First, we restrict the sample 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. Next, we include a measure of how long each model-year has been available on the market as a control, since it is well known that prices decline over the course of a model cycle, and demographic groups may vary in their propensity to purchase vehicles over their model cycles. Next, we include as a regressor the number of days that the vehicle has been on the dealer's lot—popular cars typically turn over very quickly and so dealers may be willing to reduce margins on cars that are not in high demand. Next, we include fixed effects for the city and state in which the purchaser resides, to control for fine market-level differences in demand or supply that may not be captured by state fixed-effects, especially in large states.<sup>25</sup> Finally, we also report results for two smaller, more racially homogenous states—Iowa and Wisconsin. The results, which are reported in the online Appendix, in all five cases are very similar to our main results.

#### IV(iv). *Extension to the Canadian market*

We now show that our results extend to the Canadian market. The automobile industries in Canada and the U.S. are closely integrated, and the

<sup>24</sup> Notice that these three sub-samples overlap considerably with each other. For example, the Honda Civic and the Toyota Corolla were two of the most popular cars in America in recent years, as well as both foreign-owned and compact.

<sup>25</sup> This is computationally very intensive as there are over 50,000 individual city-state combinations in the data. For this exercise we restricted the sample to the 75% of observations accounted for by larger towns and cities.



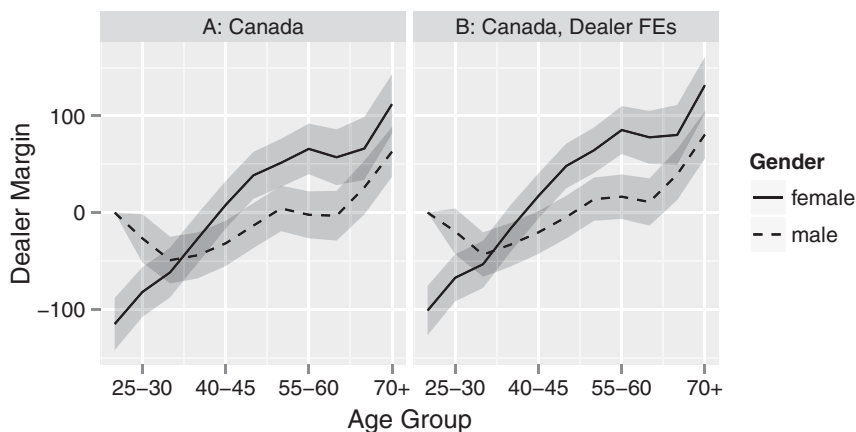


Figure 7  
Dealer Margins for Subsamples, by Gender and Age: Canadian Data

North American manufacturers, in particular, operate production on both sides of the border. All of the major domestic and foreign manufacturers offer a similar set of vehicles in both countries, with occasional differences in the names of car models, and minor differences in specifications. The process of customer negotiation for new cars is also almost identical in the U.S. and Canada.

Our data provider gave us access to a sample of approximately 1 million new car transactions in Canada. Most of the relevant variables are identical to those in the U.S. sample. One additional variable contained in the Canadian data is an identifier for the dealer at which the vehicle was purchased. As a result we can examine whether including dealer fixed-effects has any effect on the results. In the U.S. data used above, the most we could include in this regard was the geographic location of the consumer, which effectively controlled for regional variation in final prices, but did not allow for systematic differences in the behaviour of individual car dealers. This may be a concern if, for example, 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 control for this possibility by including dealer fixed-effects in the Canadian sample, which allow us to examine price differences *within* dealers.

Figure 7 presents the coefficients from running our main regression specification on the Canadian sample.<sup>26</sup> The left panel shows the main

<sup>26</sup> The full regression results are in the online Appendix.

regression while the right panel adds dealer fixed-effects. There are two main points of interest. First, the results we obtained for the United States hold broadly in Canada as well. In particular, the two main results from the U.S. sample—that there appears to be an age premium, and that this premium rises more steeply with age for women than for men—continue to hold in the Canadian market.<sup>27</sup>

The second point of interest is that adding dealer fixed-effects has virtually no effect on the estimated age and gender effects. This is visually apparent from Figure 7, and this conclusion holds up when examining the coefficients in detail. This suggests that the results in our main sample using U.S. data are not driven by systematic differences in dealer behavior.

#### IV(v). *Summary*

We emphasize three findings emerging from our results—see Figures 4–6 and 7. a) There is considerable variation in the prices paid for new cars, even after controlling for all observable features of the transaction—the difference between the 75th and 25th percentiles of the distribution of dealer margin residuals is well over a thousand dollars (compared to the average dealer margin of \$1,215). b) There appears to be a clear age premium in new car sales, with older consumers paying significantly more than younger consumers for the same vehicle, and a steady, almost monotonic, increase in margins with age. c) There also appears to be a gender divide that increases with age; as a result, older women pay the most among all consumers for a given vehicle, which is about \$240 more than the lowest paying consumers.

These results are interesting and perhaps also surprising. We control for the combination of model, state and trim of each vehicle. We also control for the timing of the transaction, the location of the buyer (through the state of residence, but also the city in a robustness check), and the dealer's cost of the purchased vehicle. We come very close to examining differences in dealer margins for identical products sold at the same time to different consumers. Therefore, any remaining differences in dealer margins beyond these controls must derive purely from idiosyncratic differences between customers. We further show that average differences across age and gender groups account for at least 20% of the remaining variation. We now turn to the question of why consumer demographics should affect the negotiating process.

<sup>27</sup> One distinction appears to be that young, female consumers in Canada outperform their male counterparts.

## V. DISCUSSION AND EXPLANATION OF THE RESULTS

In this section we consider possible explanations for our findings and discuss our results in the context of the prior literature. This section is divided into four parts. We first examine the role of competition in new car sales; we show that the demographic differences that we have identified play a more important role than competition in creating price dispersion and, moreover, that competition has no effect on our estimated age-gender premiums. Next, we discuss a number of plausible explanations for our findings that we argue can be ruled out by careful examination of the data. We then present our preferred explanation, which is consistent with additional data on changes in women's education and labor force participation. Finally, we show how our results can reconcile prior findings that appear to be in conflict.

*V(i). Competition, Demographics and Price Dispersion*

So far, we have examined dealers' ability to price discriminate among the set of buyers they encounter, and shown that there is an age gradient that is steeper for women than for men. We now examine the role of competition in this relationship. We have two goals in mind: first, to examine whether our results are in fact explained by competition among dealers.<sup>28</sup> Second, we examine whether the competitive environment changes the age-gender premiums that we have documented in the previous section.

We measure competition in three ways: the first is the number of dealers in the same county who sell vehicles of the same brand.<sup>29</sup> To do this, we obtained a separate industry dataset that lists the number of dealerships in each county that sell a particular brand of vehicle. We expect that a greater number of rival dealers would depress margins.<sup>30</sup> Our second measure of competition is simply whether the transaction takes place at the end of the month or year, when dealers are known to offer discounted prices in order to meet sales quotas. Our third measure of competition is a demand-side measure—the extent to which the vehicle is in high demand by customers, which should tend to reduce dealers' willingness to negotiate prices down.

<sup>28</sup> This could occur if certain demographic groups visit dealers facing particular competitive environments. For example, dealers in high-density urban areas usually face less competition, and may well be visited by a non-representative sample of age and gender groups. It would be interesting to examine the characteristics of buyers and the outcomes within each dealership, but our data do not include separate dealership indicators.

<sup>29</sup> Experimentation with alternative measures, including total dealers and dealers of the same manufacturer as opposed to brand, yield similar, but less precise, results.

<sup>30</sup> Note, however, that variation in dealer density is not random, so other factors correlated with dealer density could affect our estimates. For example, counties with more population will have more dealerships, so if other factors correlated with urban density determine price outcomes, or the variation in prices, then this will be picked up in our regressions.

We define high demand vehicles as models that feature the 10% lowest days to turn of vehicles sold nationally in the same quarter.

We show, in the online Appendix, that having more dealers in a county only has a small effect on prices.<sup>31</sup> The first column adds the dealer competition variables to our baseline specification (while not reporting the gender and age coefficients) and the second column estimates the regression without the demographic variables. In both cases, each additional rival lowers profits by roughly \$20 on average. Adding the demographic variables has no impact on the estimated competitive effects.

We then show that the competitive environment faced by dealers has no effect on the estimated age-gender gradients. Figure 8 divides the sample into four groups, corresponding to various numbers of rivals faced by each dealer. In all cases, the age-gender premia show the same pattern as in our full results.<sup>32</sup>

We now turn to the two other measures of competition and demand. We estimate that vehicles that are in high demand nationally, as measured by being in the lowest decile of days to turn in the same quarter, appear to command about a \$200-\$300 premium over other vehicles. However, Panel A of Figure 9 shows that high demand vehicles have no effect on the age-gender gradients. Greater willingness to sell by dealers, as measured by month-end and year-end indicators, are associated with about an \$80 discount. Again, though, there is no relationship with age-gender gradients as can be seen in Panel B of Figure 9.

Finally, we examined the relationship between competition and price dispersion, which we measure by the standard deviation of the residuals in the dealer margin regression. The online Appendix presents a figure which shows the relationship between a dealer's competitive environment and the degree of price dispersion. The figure shows no obvious effect of competition on price dispersion—although one interpretation may be a U-shaped relationship. This is consistent with the large literature on competition and price dispersion, which has found sharply opposing effects in different studies.<sup>33</sup>

<sup>31</sup> Although we estimate fixed-effects for each configuration of rival dealers, we only report the first ten coefficients. The remaining coefficients are both small and imprecisely estimated.

<sup>32</sup> Note that the magnitudes of the coefficients are not comparable across panels, since the coefficient for men under 25 is normalized to zero in each case. Also, in our transaction data, we observe the county of residence of the buyer, not the county of the seller, so our matching includes purchases where an individual purchased a car in a county with zero dealers that sell the relevant brand. These transactions are shown in Panel A.

<sup>33</sup> Borenstein and Rose [1994] found that greater competition increased price dispersion in the U.S. airline industry, while Gerardi and Shapiro [2009] found the opposite result in the very same industry. Chandra and Lederman [2016] argue that both effects are possible simultaneously at different points of the price distribution, and Dai *et al.* [2014] estimate a U-shaped relationship between competition and price dispersion that is similar to what we find, as shown in the online Appendix.

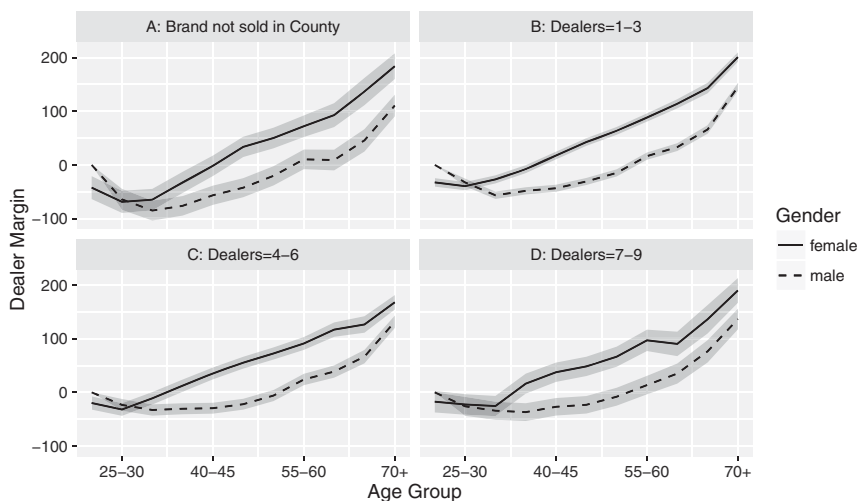


Figure 8  
Dealer Margins by Number of Rival Dealers

Our results in this subsection can be summarized as follows: we find that competition from rival dealers has only a small effect on dealer margins, far less than we estimate for average demographic characteristics of buyers. High-demand vehicles do increase dealer margins, with a magnitude similar to the extremes of our demographic variables. However, none of the three measures of competition has any effect on the age-gender gradient that we estimate in our baseline regressions.

Overall, most of the variation in dealer margins is clearly driven by idiosyncratic factors surrounding individual dealer-buyer negotiations. It is striking, though, that mean demographics can still explain around 20% of the variation and is one of the two leading explanations that we identify, along with high/low demand vehicles.

#### V(ii). *Potential Explanations for the Results*

We now consider potential explanations for our findings. We first address the concern that our results may not accurately reflect outcomes for the gender and age group associated with each transaction. In particular, one may be concerned that the person negotiating for the car is not always the same as the primary buyer listed on the invoice, and that this may be particularly likely for women and younger consumers, who may be accompanied by friends or family members negotiating on their behalf.<sup>34</sup> However, third person negotiation is unlikely to explain our pattern of results. Younger consumers—both men and women—generally obtain the best

<sup>34</sup> See Goldberg [1996] for a discussion of this issue.

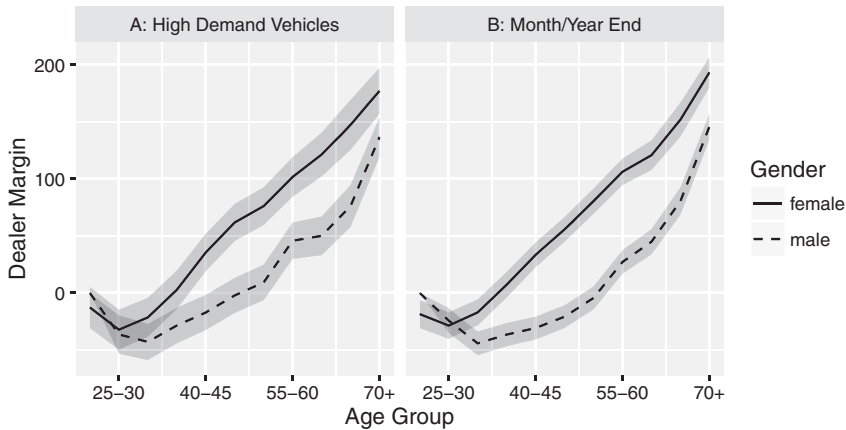



Figure 9  
Dealer Margins for Demand- and Supply-Side Factors

negotiating outcomes. If these consumers are helped in bargaining by their fathers, for example, it would be strange that these men do a better job negotiating for their children than for themselves. On the other hand, if we believe that a fraction of cars sold to women involve negotiations by male partners, our results will understate the true gender differences in negotiation. In that case, women on their own are likely to do even worse than our results indicate.


Similarly, there may be sample selection driven by marriage—women are perhaps less likely to be listed as the primary buyer on ‘family cars’ bought jointly with their husbands. If this selection removes certain women from our sample, our results will not represent the entire female population.<sup>35</sup> However, the sample selection explanation requires not just that married women are less likely to be listed as the primary buyer, but also that this proportion increases with age. In addition, it requires the sample of women purchasing cars by themselves (potentially single) to be worse negotiators, on average, than married women jointly negotiating with their partners. These possibilities cannot be ruled out, but they require unlikely conditions. We see no clear reason why single older women would inherently be worse negotiators than women of the same age who are married or in a relationship.<sup>36</sup> Further, if this were in fact true, it would imply that the results would be different for those segments that comprise a high proportion of family cars. Recall, however, that our results were no different for

<sup>35</sup> See Langer [2016] for a detailed discussion related to this issue.

<sup>36</sup> Census data indicate that unmarried women in the 40–60 age group are more likely to participate in the labor force than married women, and that the two groups have similar levels of educational attainment.

PANEL A: FEMALE CUSTOMERS  
BY AGE (%)

Source: our sample.

PANEL B: MARRIAGE RATES BY  
AGE IN 2005 (%)

Source: The American Community Survey, 2005.

Figure 10

Spousal Negotiation for Female Customers [Colour figure can be viewed at [wileyonlinelibrary.com](http://wileyonlinelibrary.com)]

SUV's—which are typical family cars—than for our full sample. We also find similar results for the Van segment, see the online Appendix.

In our dataset, the proportion of female buyers falls with customer age—see Panel A in Figure 10—partially supporting the view that older married women are less likely to be listed as the primary buyer. However, in panel B of the same figure we illustrate the likelihood of being married, by age group, for females in the United States in 2005. The probability that a woman is married crosses 60% by her mid-thirties and then begins to decline after age 45, falling quite sharply after age 60, due to the effects of both divorce and bereavement. Thus, women in their 60's or older are significantly less likely to be married than women in their 30's and 40's, suggesting that the marriage-based explanation does not influence our results.

Thus, we believe it is unlikely that our results are driven by selection—i.e., by the removal of superior negotiators from the cohort of older women in our sample. Instead, the results may be consistent with causality running in the opposite direction: the worse performance of older women in new car negotiations may lead some of them to drop out of the market altogether. If so, their main alternatives would either be to purchase used cars, or else to drive their existing vehicles for longer than average.

We do not have data on the used-car market, but we can examine the value of traded-in vehicles, which is likely to be closely correlated with their age. Our dataset contains information on the Actual Cash Value (ACV) of traded-in vehicles which is the value at which the vehicle is booked into inventory.<sup>37</sup> In Table VII we present the average trade-in values, by age category, for male and female buyers. We used data on approximately 3.5 million trade-ins, restricting the sample to those where the ACV was under \$15,000.

<sup>37</sup> The ACV is different from the amount the dealer allows the customer for the trade-in.

TABLE VII  
AVERAGE TRADE-IN VALUES (\$)

Age	Female	Male	Difference (%)
< 25	4,951	5,481	10.7
25-30	5,834	6,161	5.6
30-35	6,025	6,389	6.0
35-40	5,980	6,336	6.0
40-45	5,857	6,170	5.3
45-50	5,751	6,076	5.6
50-55	5,736	6,071	5.8
55-60	5,659	6,209	9.7
60-65	5,478	6,240	13.9
65-70	5,380	6,272	16.6
> 70	5,070	6,120	20.7

Notes: Results based on 3,545,609 transactions where a vehicle was traded in, with an ACV of under \$15,000.

The Table shows that women trade in vehicles with lower values than men; the difference is around \$300 for most age groups, but widens to about \$900 for the three oldest age groups. Proportionately, the average gender difference in trade-in values is 5–6% for consumers under age 55, but over 20% for those above 70. Clearly, older women trade in lower valued cars than younger women, relative to men of the same age. One possible explanation is that older women drive their cars for longer before trading them in, which explains, at least in part, our earlier finding that women are less represented among older cohorts of car buyers. Our results therefore suggest the possibility that we raised early in this paper—that loss aversion or perceived differences in negotiating ability may cause some demographic groups to avoid markets that involve negotiation.

Turning to other explanations, there is a large literature on gender differences in negotiations, primarily over wages, and some consensus that women perform worse in such negotiations either due to discrimination or to their own reluctance to negotiate; see Babcock and Laschever [2003], List [2004] and Leibbrandt and List [2012]. In the new car market itself, Ayres and Siegelman [1995] famously showed that women were initially offered worse terms, although these results were later contested. However, our findings argue against systematic discrimination against women, primarily because there is clear evidence that among younger cohorts women perform no worse than men, and possibly even a little better.

Analogously, there may be reasons that older consumers perform worse in negotiations. Prior research has shown that older consumers are less likely to try new brands or products; see Lambert-Pandraud and Laurent [2010]. More relevant to our study, Lambert-Pandraud *et al.* [2005] find evidence in the new car market that older customers are more likely to purchase the same brand as their existing vehicle, for reasons such as change aversion, cognitive decline, and even nostalgia.



They also consider fewer brands, fewer dealers, and fewer models than younger customers. Such an attachment to brands can allow dealers to extract higher profits from older customers. While this is a relevant explanation, it does not fully explain the gender based differences that we observe.

Another plausible explanation for our results may lie in differences in search and negotiation costs across various demographic groups. Consistent with this explanation, Morton *et al.* [2011] show that consumers with information on dealers' reservation prices capture larger shares from their margin. This is more important with access to the Internet, as online car referral services such as *Edmunds.com* often reveal dealer costs. If certain consumers are less likely to access the Internet for research, their higher search and negotiating costs could lead to poorer negotiations. Goldfarb and Prince [2008] document the digital divide in Internet use, showing that women and older consumers are less likely to use the Internet, controlling for other factors (they do not interact age and gender). We analyzed detailed data on Internet use from the Integrated Public Use Microdata Series of the Current Population Survey. While we find that older individuals in the U.S. are markedly less likely to use the Internet, the data do not indicate large differences across gender by age cohort (for details, see the online Appendix).

A related explanation for our results may be differences in income or wealth across demographic groups, which would tend to create differences in price elasticity. While it is true that women generally have lower earnings and wealth than men, our analysis of Census data suggests relatively lower wealth levels for older women, relative to men of the same age, than for younger women, which should suggest greater price sensitivity for the former.<sup>38</sup> A more detailed analysis, performed in the state-level regressions that are discussed below, shows that income and wealth differences across demographic groups do not explain our results.

#### V(iii). *Our Preferred Explanation*

Our preferred explanation for the results relates to the dramatic improvements in socio-economic outcomes for women in recent decades. Since the 1970's, women in North America have considerably narrowed their gap in education with men, and have surpassed them recently (see Goldin *et al.* [2006]). Women have also narrowed, though not closed, the wage and employment gap with men. As a result, younger women today have vastly different educational and labor market experiences than older women. Women under the age of 40 are more likely than men of the same age to have completed college, but the opposite is true for cohorts above the age

<sup>38</sup> This analysis is complicated by needing to analyze family income and wealth for married women; therefore we restricted our investigation to single men and women.

of 45 (see the relevant figure in the online Appendix, which is based on CPS data).

Female labor force participation has also grown rapidly, and as a result young women today are far more likely to be employed, relative to men of the same age, than older women when the latter were the same age. In the year 2006, almost half of women aged 25-29 were employed full time, while this was not the case for those currently aged 55-59; when this cohort was 25-29 only 37% of them were employed full time.<sup>39</sup> By contrast, in 2006, 70% of men aged 25-29 were employed full time, while those currently aged 55-59 had a similar full-time employment rate of 71% when they were the same age.<sup>40</sup>

Socio-economic trends in Canada have matched those in the U.S., with clear evidence that women have dramatically improved their levels of educational attainment and also narrowed the gap in employment with men, over the last few decades. As a result, it is the case in Canada, too, that older women today have lower levels of education and labor force participation, relative to men, than women in their 20's. Therefore, these demographic trends may well explain the similar pattern of results that we established for the Canadian market.

We conjecture that similar educational attainment and labor force participation of young women, relative to men of the same age, allows them to perform as well in negotiations. Women above the age of 60 are much less likely to have completed high-school or college, or to have been employed full-time when they were younger. These differences can potentially cause older women to have lower information in the new car market and perhaps also to negotiate with less confidence. This is likely to play an important role given the complexity of negotiations, which include discussions around financing, monthly payments, trade-in allowances and service contracts. Correspondingly, this may explain why younger women of today are as good negotiators as their male counterparts. Note that if this explanation is correct, it is unlikely that today's cohort of younger women will do worse than their male counterparts as they age. The difference in relationships observed is therefore likely to be specific to the cohorts currently observed, rather than reflecting an ongoing effect.<sup>41</sup>

<sup>39</sup> We are comparing employment outcomes of those born between year 1976-80, with those born between 1946-50. We compare the likelihood of working at a fixed age for each cohort, because women's labor force participation declines with age in every cohort. Women are most likely to be employed when aged 25-29, which is why we chose this age group as our comparison point. However, the trends are very similar if we instead compare labor force participation at other ages.

<sup>40</sup> Data on workforce participation are from the Integrated Public Use Microdata Series of the Current Population Survey.

<sup>41</sup> We also performed a state-by-state analysis to examine whether these trends were particularly apparent in U.S. states that had the greatest or lowest socio-economic changes for women. These results were inconclusive, for two main reasons. First, we assigned state-wide

We also see evidence that transactions with lower levels of complexity have less dispersed prices and smaller age and gender-related premiums. Recall, from Section IV., that the smallest age-gender gradient was obtained on the sub-sample of the data where customers did not trade in an existing vehicle. Negotiations over the trade-in constitute an important part of the overall car buying process, along with discussions about other issues. Dealers, who perform many such negotiations, have considerably more experience and information on these matters than consumers, and it is likely that they can use the complexity of transactions to their advantage. For example, they may offer consumers seemingly attractive terms on the new car, only to make up the difference through a lower trade-in allowance, higher interest rates, profitable service contracts, extended warranties, and other fees. It is revealing, therefore, that transactions that do not feature trade-ins also have smaller differences across demographic groups, suggesting that trade-in allowances are one way for dealers to profit on less informed customers. This fits well with the larger explanation that consumers with more information, which could be driven by better education or even simply by better access to information on the Internet, perform better in price negotiations.

V(iv). *Reconciling Earlier Findings*

We now discuss our findings in light of the prior literature in order to reconcile the conflicting results in these studies. We show that our results encompass the seemingly opposing results in some of these studies, and that our focus on the age of customers is particularly helpful in understanding these issues. Moreover, our access to a much larger sample, with more detailed information on vehicle characteristics, appears to explain important differences among prior studies.

Ayres and Siegelman [1995] used an audit study to show that women performed worse than men in the new car market, in terms of both initial offers by dealers and final negotiated prices. Following this, Goldberg [1996] and Harless and Hoffer [2002], using survey data, found no evidence of a gender gap in final transaction prices.<sup>42</sup> It is important to note that none of these

averages for educational attainment and labor force participation to the set of car buyers in each state, as we lack microdata on these variables for buyers in our sample. But this is a crude measure since the customers in our sample form a small—and likely unrepresentative—fraction of the population in each state and year. Second, the only way to assess labor force participation of older cohorts when they were younger is by assigning to them historical state-wide averages for the corresponding periods. However, this is severely complicated by the high mobility of Americans; according to Census estimates, only 58% of U.S. adults in 2010 lived in the same state in which they were born.

<sup>42</sup> However, Harless and Hoffer [2002] did find some evidence that older consumers pay more in negotiations, a finding that was recently repeated in the French automobile industry by D'Haultfoeuille *et al.* [2014].

studies examined the joint distribution of age and gender and therefore their results were focused entirely on *average* differences in gender.

Focusing on gender as a purely binary identity can obscure important variation within each gender. Indeed, Goldberg [1996] speculated that one cause for the discrepancy between her study and that of Ayres and Siegelman [1995] may have been potentially higher variation in reservation prices for female buyers. Our results are consistent with Goldberg's conjecture: our main results—in Table VI and Figure 4—show that the difference in dealer margins between the highest and lowest paying women is considerably higher than for the corresponding groups of men.

Importantly, our results encompass both sets of prior results; either result regarding the effect of gender can be found, depending on the age of the sample. Examining only the younger cohort of buyers—those aged under 35 or so—would lead to the conclusion that there is no difference in transaction prices between genders, which would be consistent with both Goldberg [1996] and Harless and Hoffer [2002]. However, focusing on the older cohorts—buyers above age 50 or so—would lead to the conclusion that women fare worse than men in new car negotiations, which would be entirely consistent with Ayres and Siegelman [1995].

Furthermore, when we recreate the specific regression environment of earlier studies we can better understand, and reconcile, their conclusions. Goldberg [1996] found that women do not fare worse in new car negotiations than men, in contrast to Ayres and Siegelman [1995]. A limitation of Goldberg's study was the lack of detailed data on vehicle trims and options. As we showed in Table V, this information is directly correlated with gender: among the buyers of a given model, women choose systematically cheaper trims and options than do men. Omitting this information would therefore tend to bias the results towards showing that women fare better in negotiations than they in fact do.

Indeed, when we repeat our regressions without using our knowledge of a vehicle's trim and options—i.e., using the transaction price as the dependent variable rather than the dealer margin, which embeds aftermarket options, and without including the vehicle trim in the definition of fixed-effects—we find that women appear to negotiate better prices than do men.<sup>43</sup> Thus, it seems likely that Goldberg [1996] underestimated the prices paid by female buyers, and that at least part of the difference between Goldberg [1996] and Ayres and Siegelman [1995] is due to the lack of information in the former on the different choices made by women with respect to vehicle trim and options.

Harless and Hoffer [2002] used a similar measure of dealer margin in their study of gender, but they used customer age as a control, rather than interacting it with gender as we do. When we repeat their regression

<sup>43</sup> The difference in the gender premium between this specification and a similar regression incorporating all the trim and option data is around \$160.

specification—i.e., without interacting age and gender—we find a similar result to theirs: women *on average* do not have different outcomes from men.<sup>44</sup> However, this obscures the important gender differences across age cohorts that we identify in our full results.

While comparing these results, it is important to keep in mind the very different time period in our paper compared to earlier studies, especially Goldberg [1996] and Ayres and Siegelman [1995], who use data from the 1980's. Since that time, a number of factors are likely to have changed the new car negotiation process, such as consumers' access to the Internet, and the technology and information available to dealers. Nevertheless, we believe that our emphasis on consumer age cohorts is an important explanation of the different results. In fact, the consumers who constituted the samples in those studies would have shown up as older consumers in our own data which are from 20 years later. Thus, our finding of large gender differences in older cohorts appears fully consistent with the average gender differences found in Ayres and Siegelman [1995], as well as with Goldberg [1996], given that we argue that Goldberg's result of small gender differences can at least partly be attributed to the low level of detail on vehicle characteristics in her study.

Moreover, our focus on the joint distribution of age and gender is just one possible way in which the conflicting prior results can be reconciled. More generally, our findings suggest that there is likely to be considerable heterogeneity in negotiating ability within each gender. Some of this may be purely idiosyncratic, i.e., particular to individuals. However, to the extent that some of this heterogeneity is systematic, it appears to be correlated with consumer age, but age may just be a proxy for some other characteristic that we have been unable to fully identify.

In this regard, an interesting parallel to our paper is Langer [2016], who shows that preferences and demand elasticities for new cars vary considerably by marital status—which is of course highly correlated with age. While Langer argues that firms do not exploit these differences to fully price discriminate, our results complement her finding that single people are more price sensitive than married people, given that we find younger consumers paying less on average.<sup>45</sup> We view our results as extending those in Langer [2016] since we can break down the gender-based results by the age of consumers, in a manner that has not been done previously in the literature.

<sup>44</sup> This regression implies a very small female premium of \$11, which is not statistically significant.

<sup>45</sup> There are a number of important differences between this paper and Langer [2016]: she uses a structural model to predict optimal markups from transaction price data, whereas we use reduced-form techniques based on direct calculation of dealer margins.

## VI. CONCLUSIONS

In this paper we examine price negotiations in the new car market, with an emphasis on how age and gender characteristics relate to disparate outcomes. We began by establishing large variation in prices paid for almost identical new cars. We then demonstrate systematic differences in how consumers of each gender, and various age groups, perform in price negotiations. In general, older consumers pay more for new cars, but the trend is particularly stark among older women. As a result, women above the age of 70 generate almost \$250 more in dealer margins—which is about 20% of the average dealer margin—than the lowest paying customers, even after controlling for all observable aspects of the transaction. It is revealing, therefore, that older women are also the least represented among new car buyers. Our results are robust to cutting the data in many ways, and to adding a large number of covariates.

We also see evidence that younger women do as well, or better, than men of their age in the new vehicle market. This is concurrent with the reversal of the gender gap in education, and the narrowing of the gender gap in employment and wages. It could be that the rapid increase in women's education, as well as the improvement in their earnings and work experience relative to men, has given women better information and more confidence while conducting price negotiations. Therefore it could also be that the improvements women have seen from more advanced education and greater work opportunities are not restricted only to the labor market. This has important implications for other markets involving negotiations, such as the housing market.

The main limitation of drawing this conclusion is that our results are based on a cross-section of transactions, albeit one that employs a very large sample of consumers drawn from the entire range of the age distribution. Confirmation of our hypothesis that gender differences in negotiation are related to educational and labor force outcomes requires a panel that tracks cohorts over a much longer period and uses micro data on socio-economic characteristics. Since such data do not currently exist in the automobile industry, our findings should be regarded as suggestive until they can be confirmed in future research.

Irrespective of the reasons underlying our findings, further research on markets with negotiated prices is important. The existence of a negotiated markup in the automobile sector is indicative of market power. One potential solution for reducing age and gender based disadvantages is to address the source of this market power. If one of the reasons for disparate outcomes is the lack of information, increasing the flow of information and ease of its access will help. Establishing the existence of demographic-based disparities in other negotiated price markets, and identifying reasons behind these, should be goals for future research in this area.

## REFERENCES

- Allen, J.; Clark, R. and Houde, J.-F., 2014, 'Search Frictions and Market Power in Negotiated Price Markets,' technical report (National Bureau of Economic Research, Cambridge, Massachusetts, U.S.A.).
- Ayres, I. and Siegelman, P., 1995, 'Race and Gender Discrimination in Bargaining for a New Car,' *American Economic Review*, 85(3), pp. 304–321.
- Babcock, L. R. and Laschever, S., 2003, *Women Don't Ask: Negotiation and the Gender Divide* (Princeton University Press, Princeton, New Jersey, U.S.A.).
- Bester, H., 1993, 'Bargaining versus Price Competition in Markets with Quality Uncertainty,' *The American Economic Review*, pp. 278–288.
- Borenstein, S. and Rose, N. L., 1994, 'Competition and Price Dispersion in the U.S. Airline Industry,' *Journal of Political Economy*, pp. 653–683.
- Busse, M. R.; Israili, A. and Zettelmeyer, F., 2013, 'Repairing the Damage: The Effect of Price Expectations on Auto-Repair Price Quotes,' working paper 19154. (National Bureau of Economic Research, Cambridge, Massachusetts, U.S.A.)
- Busse, M. R. and Silva-Risso, J. M., 2010, 'One Discriminatory Rent' or 'Double Jeopardy': Multicomponent Negotiation for New Car Purchases,' *The American Economic Review*, pp. 470–474.
- Chandra, A. and Lederman, M., 2016, 'Revisiting the Relationship between Competition and Price Discrimination,' working paper, (Rotman School of Management, University of Toronto, Ontario, Canada).
- Chandra, A. and Tappata, M., 2011, 'Consumer Search and Dynamic Price Dispersion: An Application to Gasoline Markets,' *The RAND Journal of Economics*, 42(4), pp. 681–704.
- Chen, Y.; Yang, S. and Zhao, Y., 2008, 'A Simultaneous Model of Consumer Brand Choice and Negotiated Price,' *Management Science*, 54(3), pp. 538–549.
- Dai, M.; Liu, Q. and Serfes, K., 2014, 'Is the Effect of Competition on Price Dispersion Nonmonotonic? Evidence from the U.S. Airline Industry,' *Review of Economics and Statistics*, 96(1), pp. 161–170.
- D'Haultfoeuille, X.; Durrmeyer, I. and Février, P., 2014, 'Automobile Prices in Market Equilibrium with Unobserved Price Discrimination,' technical report (Free University of Berlin, Humboldt University of Berlin, University of Bonn, University of Mannheim, University of Munich, Germany).
- Gerardi, K. S. and Shapiro, A. H., 2009, 'Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry,' *Journal of Political Economy*, 117(1), pp. 1–37.
- Goldberg, P. K., 1996, 'Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey,' *Journal of Political Economy*, 104(3), pp. 622–64.
- Goldfarb, A. and Prince, J., 2008, 'Internet Adoption and Usage Patterns are Different: Implications for the Digital Divide,' *Information Economics and Policy*, 20(1), pp. 2–15.
- Goldin, C.; Katz, L. F. and Kuziemko, I., 2006, 'The Homecoming of American College Women: The Reversal of the College Gender Gap,' *Journal of Economic Perspectives*, 20(4), pp. 133–156.
- Harding, J. P.; Rosenthal, S. S. and Sirmans, C., 2003, 'Estimating Bargaining Power in the Market for Existing Homes,' *Review of Economics and Statistics*, 85(1), pp. 178–188.
- Harless, D. W. and Hoffer, G. E., 2002, 'Do Women Pay More for New Vehicles? Evidence from Transaction Price Data,' *American Economic Review*, 92(1), pp. 270–279.

- Jindal, P. and Newberry, P. W., 2015, 'Price Negotiations and Bargaining Costs,' Ssrn Working Paper (Available at SSRN 2471664).
- Lambert-Pandraud, R. and Laurent, G., 2010, 'Why do Older Consumers Buy Older Brands? The Role of Attachment and Declining Innovativeness,' *Journal of Marketing*, 74(5), pp. 104–121.
- Lambert-Pandraud, R.; Laurent, G. and Lapersonne, E., 2005, 'Repeat Purchasing of New Automobiles by Older Consumers: Empirical Evidence and Interpretations,' *Journal of Marketing*, 69(2), pp. 97–113.
- Langer, A., 2016, '(Dis)Incentives for Demographic Price Discrimination in the New Vehicle Market,' manuscript (University of Michigan, Ann Arbor, Michigan, U.S.A.).
- Leibbrandt, A. and List, J. A., 2012, 'Do Women Avoid Salary Negotiations? Evidence from a Large Scale Natural Field Experiment,' technical report (National Bureau of Economic Research, Cambridge, Massachusetts, U.S.A.).
- List, J. A., 2004, 'The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field,' *The Quarterly Journal of Economics*, 119(1), pp. 49–89.
- Morton, F. S.; Silva-Risso, J. and Zettelmeyer, F., 2011, 'What Matters in a Price Negotiation: Evidence from the U.S. Auto Retailing Industry,' *Quantitative Marketing and Economics*, 9(4), pp. 365–402.
- Morton, F. S.; Zettelmeyer, F. and Silva-Risso, J., 2001, 'Internet Car Retailing,' *The Journal of Industrial Economics*, 49(4), pp. 501–519.
- Morton, F. S.; Zettelmeyer, F. and Silva-Risso, J., 2003, 'Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities?' *Quantitative Marketing and Economics*, 1(1), pp. 65–92.
- Sallee, J. M., 2011, 'The Surprising Incidence of Tax Credits for the Toyota Prius,' *American Economic Journal: Economic Policy*, pp. 189–219.
- Stuhlmacher, A. F. and Walters, A. E., 1999, 'Gender Differences in Negotiation Outcome: A Meta-Analysis,' *Personnel Psychology*, 52(3), pp. 653–677.
- Wang, R., 1995, 'Bargaining versus Posted-Price Selling,' *European Economic Review*, 39(9), pp. 1747–1764.
- Zhu, R.; Chen, X. and Dasgupta, S., 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), pp. 159–170.

#### SUPPORTING INFORMATION

Additional supporting information may be found in the online version of this article at the publisher's web-site: <http://wileyonlinelibrary.com/journal/joie> or via the Journal's website: <http://www.jindec.org>