

Posted Price and Haggling in the Used Car Market*

Guofang Huang[†]

JOB MARKET PAPER

Abstract

Though haggling has been the conventional way for auto retailers to sell cars, the last two decades have witnessed the systematic adoption of no-haggle prices by many large dealerships, including the largest used car dealership, Carmax. This paper develops a structural empirical model to estimate dealers' profits under posted price and haggling and investigate the conditions leading to the heterogenous choices of pricing policies that we observe. The model incorporates a simple class of bargaining mechanisms into the standard BLP model. With the extension, the paper is able to estimate the product-level demand system using data with only list prices and estimate the discounts achieved in haggling. The counterfactual experiments based on the estimates yield a few interesting findings. First, dealers using posted price would see their profits increase only slightly if they haggled, whereas the haggling dealers' profits would drop significantly if they switched to posting prices. Second, the market power enjoyed by the non-haggling dealerships through vertical differentiation and carrying large number of models is a major factor contributing to the variation in the relative performance of posted price and haggling.

JEL Classifications: L11, L13

Key Words: Haggling, Posted price, Competition, Price Discrimination

*I'm deeply indebted to my committee members, Hülya Eraslan, Joe Harrington, Matt Shum, and Tiemen Woutersen for their persistent help and encouragement throughout the process of my dissertation. I also thank Colleen Carey, Yingyao Hu, Maia Linask, Ariel Pakes, and Marc Remer for their helpful comments.

[†]Department of Economics, Johns Hopkins University. Email: ghuang15@jhu.edu. Website: <http://econ.jhu.edu/jobmarket/2009/HuangG>. Comments are welcome.

1 Introduction

In the auto retail market negotiation is the traditional way to reach a transaction price. The conventional wisdom in this market is that negotiation allows dealers to reap higher profit from consumers not interested in or not good at negotiation and, at the same time, get more business from the price-sensitive buyers. Yet, starting in the early 1990s, some dealerships began to adopt haggle-free posted prices. For example, Carmax, founded by Circuit City in 1993, sells mostly used cars at no-haggle prices. AutoNation, the largest new car retailer in the US, started experimenting with no-haggle pricing in its Denver stores in early 1999, and has recently adopted no-haggle pricing systematically throughout its stores in the US. Many other large auto dealership chains have also converted to no-haggle pricing.

The questions I pose in this paper are why some dealerships choose to haggle while others do not, and what are the market conditions that lead a dealership to choose one policy versus the other. I address the questions in two steps: first, I ask how do dealers' current profits compare to those obtainable under alternative pricing policies. In the second step, I ask what are the conditions leading a dealership to choose one pricing policy versus the other. I attack the questions using structural empirical methods.

My general strategy is to estimate a demand system, and use it to simulate sellers' profits under various pricing policies. Conclusions will be based on the comparisons of the various profits. The strategy relies on two key steps. The first one is the ability to estimate the actual profits for all dealerships. The second one is a consistently estimated demand system, from which the profits under alternative pricing policies can be simulated. The main difficulty with the implementation of this approach is that I only observe the list prices for cars and do not observe the actual transaction prices.¹ A direct application of the random-coefficient discrete-choice model developed by Berry, Levinsohn, and Pakes (1995) (hereafter referred to as BLP) would lead to inconsistent estimates, especially of the coefficient of price and actual profits. To get around this difficulty, I incorporate actual pricing policies into the standard BLP model, which enables me to consistently estimate the demand system *and* recover the haggling dealers' unobserved pricing rules and thus the actual profits. More precisely, I assume that each haggling dealership uses a screening strategy consisting of two offers, (p, \tilde{p}) , sequentially presented to buyers, and that the cost to buyers of waiting for the second offer is measured by a cost-to-negotiation factor, $\beta \in (0, 1)$. I also assume that only the seller makes offers in the haggling. An interested buyer with a reservation value for a car would self-select into either taking the initial offer or waiting for the discount. So, the transaction prices would be determined endogenously within the model. The standard BLP model becomes a special case of the augmented model with the discounts offered by all sellers being zero.

Following BLP and Petrin (2002), I use the General Method of Moments (GMM) to estimate the

¹It has been pointed out by several authors, such as Berry, Levinsohn and Pakes (1995) and Goldberg (1995), that the product prices observed by researchers are not necessarily the prices that are directly affecting market players' decisions. The lack of actual transaction prices is common for the empirical research of the auto market and many intermediate good markets.

demand system and the discounts achieved in the haggling. Because I have non-haggling as well as haggling dealers in my data, I am able to identify in the augmented model the unobserved discounts through matching the model-predicted and the observed product market shares.² The estimation algorithm in BLP can be modified to work for the augmented model. The key contraction mapping that BLP used to invert the market share functions is still a contraction mapping in the augmented model. My estimates of the unobserved discounts are the same in magnitude as those found in previous experiments as well as in the actual buying experiences at the haggling dealerships. For the medium-size dealerships in my data, the estimates of their price-cost margins are very close to the market average price-cost margins that I derived using external sources of information.

For the empirical investigation of this paper, I treat the observed market with the formats of posted price and haggling as an asymmetric equilibrium in pricing policies. For computing and comparing the profits under various pricing policies, I focus on the scenarios that exactly match the defining conditions of a Nash Equilibrium in pricing policies. That is, I first estimate the actual profits for every dealership, and then compute for each dealership the profits under alternative pricing strategies while fixing the pricing strategies of all the other dealerships. Within each dealership, I compare the estimated profits under the actual pricing policies with the simulated profits under alternative pricing policies to reveal the relative performance of the two popular pricing formats. To further investigate the relationship between competition and the choice of pricing policies, I create two sets of counterfactual dealerships. The dealerships in the first set differ in the car qualities they supply, and those in the second set sell a different number of models. I employ these two sources of variation in the competition faced by the dealerships to see how competition may affect their choices of pricing policies.

My research finds that, first, the relative performance of posted price and haggling varies across dealerships. The dealerships using posted prices would only see their profits increased slightly if they haggled, whereas the haggling dealerships' profits would drop significantly if they switched to posting prices. Second, I find that market power derived from two sources has significant impact on dealerships' choices of pricing policies. The two sources of market power are respectively the vertical differentiation in qualities and the variety of models carried by a dealership. I find that sellers selling used cars with rare qualities and/or carrying a large number of models are more likely to use posted price.

One additional point worth mentioning about estimating the demand system is that using list prices as the transaction prices would lead to underestimates of price coefficients and price elasticities. I compare the own price elasticities of demand estimated under the augmented model and under the model treating list prices as transaction prices. The comparison shows that simply treating all list prices as the transaction prices significantly underestimates the own price elasticities of demand. The basic underlying logic is intuitive. As the transaction prices are only likely to be lower than the list prices for haggling dealers, linking the list prices to the observed market shares

²I observe the identity of the haggling dealers but not their exact pricing rules. I also have the list prices and the detailed characteristics of every used car sold by the dealers. The identification will be shown in detail later.

underestimates the negative impact of price on market shares.

My research belongs to the literature that studies the interaction between competition and price discrimination. My paper is most closely related to the paper of Borenstein (1985). In that paper, using numerical methods, Borenstein employs Salop's (1979) circular product differentiation model with single-product firms to study the price discrimination under free entry. He numerically computes and compares the *symmetric* equilibria under three scenarios: the price-discrimination-prohibited case, third-degree price discrimination and a simple form of second-degree price discrimination. His comparison results show that competition under free-entry does *not eliminate* price discrimination; price discrimination based on reservation values (vertical preference) on average increases firms' profits, total output and consumers' surplus; price discrimination based on brand preference (horizontal preference) on average increases firms' profit and output by relatively smaller amounts, and slightly decreases consumers' surplus. He concludes there is no evidence that competition can prevent firms from price discrimination. In comparison, my results go further by suggesting that *competition may actually encourage price discrimination*.

Riley and Zeckhauser (1983) study the optimal pricing policy for a monopolist in a dynamic environment. They find that quoting take-it-or-leave-it prices to buyers sequentially coming into the market is the monopolist's optimal pricing policy. Their result is robust to sellers having to gradually learn about the distribution of buyers' reservation prices, the number of objects for sale, and the size of the buyer population. One of my empirical findings is that posting price is a better choice for sellers with larger market power. So, in some sense the empirical finding is consistent with the theoretical result of the extreme case of monopoly pricing.

Other papers on second-degree price discrimination and imperfect competition usually study sellers' equilibrium nonlinear pricing schemes for a product with a continuum of qualities.³ In contrast the goal of this paper is to find out sellers' optimal pricing schemes for some given products when facing consumers with unobserved heterogeneous preferences. Although the basic idea of price discrimination is the same, the mathematical formulations of the two mechanism design problems are slightly different. For example, the cost function is commonly assumed to be increasing and convex in quality in the nonlinear pricing problem. Whereas for the optimal pricing problem, the marginal cost of a product is constant, and moreover it often involves additional screening costs to the seller to sell to consumers with low reservation values. I am not aware of any theoretical article that directly studies the interaction between competition and optimal pricing strategies.

³Stole (2007) reviews the theoretical literature of price discrimination in markets with imperfect competition. Results of theoretical models of second-degree price discrimination generally depend on the specific assumptions that authors make. To make solving the equilibrium of the price discrimination game easier, the results are almost always obtained for symmetric equilibria, and in many cases for duopoly markets. For example, Armstrong and Vickers (2001) and Rochet and Stole (2002) are two theoretical papers that analyze price discrimination in rich environments similar to the market we're interested in. Both papers are aimed at finding the form of equilibrium nonlinear pricing schemes when consumers differ in both their vertical and horizontal preferences, and both also focus on analyzing the symmetric duopoly case under the assumption that the market is fully covered. Their results show that sellers will price discriminate in the symmetric equilibria. Yang and Ye (2008) study the interaction between competition and price discrimination without the assumption of the market being fully covered. They find that in comparison to the monopoly case, duopoly involves less quality distortion and more sales to consumers. For empirical papers on nonlinear pricing, see, for example, Leslie (2004) and Mcmanus (2007).

This paper has several contributions to the literature. First, it contributes to the literature of price discrimination under imperfect competition by empirically revealing some important factors affecting the choice of pricing rules in an *asymmetric* equilibrium in pricing policies⁴. My paper also contributes to the empirical Industrial Organization (IO) literature on demand estimation. McFadden (1978, 1981), Berry (1994), Berry, Levinsohn, Pakes (1995), Petrin (2002), Berry, Levinsohn, Pakes (2004) among others have developed the now widely-applied discrete-choice demand model with random coefficients. My exercise shows that the framework could be extended by incorporating actual pricing institutions into the model. Such an extension would especially improve the estimation of demand elasticities, and consequently be important for applications that rely critically on accurately estimated price elasticities.

The rest of the paper is organized as follows. In section two I introduce the industry background of the used car market. I set up my empirical model in the third section. Section four is the introduction of my data. In section five I explain the identification of my model and the estimation procedure. Section six presents and discusses the estimation results. Section seven presents the counterfactual experiments. The last section concludes.

2 The Used Car Retail Market

The total sales of the U.S. used vehicle market was about \$340 billion in 2007⁵. In the same year there were an estimated 41.4 million used vehicles sold in the U.S., compared to 16.2 million new vehicles sold in the US. These numbers make the used car market the largest retail segment in the U.S by revenue⁶.

The used car market is highly competitive. It includes approximately 21,800 franchised new car dealers, 44,000 independent dealers, as well as millions of private individuals. The major players in the market are CarMax (the largest specialized used car retailer), and large franchised new car dealers, such as AutoNation and Penske. These firms sell mostly late model vehicles that are 1 to 6 years old. Independent dealers sell predominantly older cars with higher mileage.

Dealers acquire their used car inventory through several sources. Franchised new car dealers get a large portion of their used car inventory through buyer trade-in. Dealers also acquire used cars through local and regional auctions, wholesalers, leasing companies and rental companies. Overall, acquiring used cars directly from consumers, i.e. through trade-ins, seems to be the dealers' favorite source. The competition between used car dealers also exists in the acquisition market, which tends to drive up dealers' acquisition cost. Dealers usually finance their inventory through internal funds and short-term loans from banks.

Used car dealers typically sell their cars in one of two ways: they either put up list prices (sticker prices on windshields and prices online) and let their sales team negotiate with buyers to

⁴Our findings are specifically based on the auto retail market. So, it is also of value to policy makers and business managers interested in analyzing the market.

⁵The total revenue of the used car market is about the same as that of the new car market.

⁶These numbers are cited from the 2007 Carmax Annual Report.

reach transaction prices; or they only sell at list prices, allowing no haggling at all. Negotiation is the traditional and the more common way of selling cars. The often-discussed problem with the haggling method is that many customers feel uncomfortable about haggling with the salesperson. Starting in 1990, GM's Saturn was the first to sell cars at no-haggle prices. CarMax, established in 1993 and focused on used car sales, is another pioneer in selling cars using the haggle-free format.⁷ Many large franchised new car dealers have also been experimenting with the no-haggle way of selling cars. For example, AutoNation, the largest auto retailer in the U.S., experimented with no-haggle prices, together with the superstore concept, first in its dealerships in Denver, Colorado in early 1999. The program was at first quite successful: both its market share and profit in Denver increased significantly in that year. Later that year, AutoNation switched to selling cars at no-haggle prices in its dealerships in Tampa, Florida.⁸ In 2006, AutoNation officially committed to no-haggle pricing by adopting the "Smart Choice" system in its dealerships nationwide. More recently, Lithia Motors Inc., the 8th largest new car retailer in the U.S., switched to no-haggle pricing in 2007. For auto manufacturers, Ford and GM have also been picking out some of their models and asking them to be sold by their dealers under no-haggle prices in selected locations. On the other hand, I do not observe any independent dealers formally committed to no-haggle pricing.

The choice of pricing policy is a major decision for the auto dealers. The implementation of a new pricing policy involves changes ranging from the compensation structure of salespersons⁹ to the management information system¹⁰, and it may make a big difference in dealers' profits. Although a majority of the auto dealerships still sell their cars through negotiating with buyers, the switch towards posting price is surprising given the conventional wisdom in the industry. How much would dealers' profits change if they used alternative pricing strategies? What causes dealerships in the same market to apply different pricing policies? Answers to these questions would help us better understand the conditions that have shaped the sellers' choices of pricing strategies in this important market. In the following, I use novel dealership-level aggregate sales data from the used car market to address these questions through structural empirical methods¹¹.

⁷The no-haggle format initially raised some anti-trust concerns. In 1994 the US Justice Department looked into the no-haggle format of selling cars, citing the anti-trust laws. The details can be found in a 1994 report, "U.S. Investigating No-Haggle Car Pricing", by the New York Times.

⁸The dealerships in the two places were packaged under a new brand AutoNation USA. But the initiative was abandoned by the end of 1999 because "the superstore chains invested too much in real estate, did not have enough late-model used cars and set prices too high". The main problem with the experiment was not the effectiveness of no-haggle pricing, but the "superstore" business model. The previous quote is from the article "Retail misadventure struck a nerve with dealers" from www.autonews.com.

⁹The compensations of salespersons at the bargaining dealership are often dependent on the profit margin on the deals completed by them, whereas salespersons are compensated based on a fixed dollar-per-vehicle standard at fixed pricing dealerships.

¹⁰A management information system normally collects demand information, monitors sales trend, generates pricing adjustment recommendations and optimizes inventory turnover rate.

¹¹Alternatively two other approaches may be taken to answer my research questions. One is to solve the equilibria in pricing strategies for a similar environment with imperfect competition. First I note that for the market I am interested in, an appropriate imperfect competition model include more than two multi-product firms, products with multiple characteristics, and consumers with privately known multi-dimensional (both vertical and horizontal) preferences. Taking the theoretical approach for my investigation is very difficult because for example it entails solving asymmetric equilibria in pricing strategies for multi-product firms. To investigate the questions through

3 Empirical Model

3.1 General Description

I follow the framework developed by Berry, Levinsohn, Pakes (1995) in setting up my demand model. Their original model can not be applied directly to investigate the choice of pricing policies as the sellers in their model are implicitly assumed to always offer no-haggle prices. Hereafter, I will call the direct application of the BLP framework the “baseline model”. For the purpose of this paper, I extend the baseline model by incorporating sellers’ actual pricing policies into the model. As I explain in more detail later, when estimating the demand system this extension also enables me to solve the data problem of not observing the actual transaction prices for the haggling dealerships.

A few issues naturally arise when modeling the used car retail market. The first such issue is about the concept of products. The problem here is twofold. One is that I need to model the competition between dealers, and not just between brands as in the new car market; and the other is that there are too many products to handle in the empirical model because used cars of the same model differ in their age and accumulated mileage. In order to model the competition between dealerships, I define a product by “model/dealership”: the same model sold by different used car dealerships are treated as different products. Several dealerships may carry identical models, but they may still differ in the year and mileage they focus on. To make the number of products manageable, I aggregate the used cars within each dealership to the model level and treat the average mileage and age of a model sold by a dealership as additional product characteristics. And accordingly, I take the average list price of each product as its list price. A natural concern about the approach is that the aggregation may cause bias in my estimates. I follow the suggestion of McFadden (1978) to deal with the problem by adding, as additional product characteristics, the log of each model’s average inventory size, and the variances and covariances of age, mileage, and price.¹²

The second issue is that the price I observe for each car is the list price instead of the transaction price. The observed list price is indeed the transaction price at the non-haggling dealerships, but it is often not the transaction price at the haggling dealerships. Similar data problems also exist in the new car retail market and other retail markets where the transaction prices may differ from the observed posted price when some consumers receive discounts or use coupons. The problem

reduced-form empirical analysis, one might want to compare the ex ante and ex post data in a market in which some dealers switched pricing strategies. But data at the dealership level are rare, and it would be even harder to find data covering a period in which some dealers switched pricing strategy. In addition, though we may know which dealers are haggling, we normally do not observe the detailed pricing strategies used by the dealerships, which makes it difficult to measure the dealers’ actual profits. To get around these difficulties, I take the structural empirical approach by taking advantage of the unique dealership-level aggregate sales data I collected.

¹²McFadden (1978)’s problem when studying housing choice is that there are too many products if one treats each individual house as a different product. Assuming that the house characteristics follow a log normal distribution, he shows that by including the log of the number of houses of each type of house, and the variances and covariances of house characteristics for each type of house, aggregating individual houses up to the level of type of houses does not cause bias in the estimation.

with directly using list prices as transaction prices in the estimation is that the discounts off the list prices are always nonnegative and are not offered randomly. Specifically for the auto retail market, the buyers normally negotiate with the dealers to reach transaction prices. So, the transaction price is bounded above by the list price and is strictly lower than the list price for a selected group of consumers. Hence using the observed list prices directly as transaction prices could lead to an inconsistent estimate of the demand system. I get around this data problem by incorporating dealers' actual pricing policies into the demand model. To simulate dealers' profits under alternative pricing policies I also need a model that can accommodate various pricing policies. Explicitly modeling dealers' pricing policies will serve both purposes.

I describe here the simple model I use for the negotiation process, and will discuss the motivations of using the model in the next subsection. I take the dealer's main goal in the negotiation as to screen consumers with unobserved reservation values for price discrimination. Thus I assume that in the negotiation the seller sequentially presents two offers. For my application, the observed list price p will be taken as the initial offer price. The dealer waits for some fixed time before presenting a second lower price \tilde{p} , which is not observed by the econometrician. The delay in presenting the second offer has a negative impact on the buyer's utility: the buyer's net surplus from purchasing a car at the second price will be discounted by a factor $\beta \in (0, 1)$.¹³ An interested buyer either accepts the initial offer or waits for the dealer to offer a lower price, but she cannot present a counteroffer. I assume that each haggling dealership uses the same strategy, i.e. the same discount, to sell all its cars. I will allow the discounts to vary when checking the empirical relevance of such flexibility. Let d be the discount reflected in the dealer's final offer price, i.e. $d = p - \tilde{p}$. Then, given the observed list price p , (β, d) is a complete description of a dealer's screening strategy.¹⁴ The identification of (β, d) will allow me to estimate the demand system by using only list prices, and with the recovered discounts, I can estimate the actual profits obtained by the haggling dealerships. It is worth emphasizing that in my estimation I am not assuming that the sellers' screening strategies are necessarily their optimal strategies.

Lastly, the choice of buying a new car cannot be ignored. As my focus is on the used car market, I treat buying new cars as a single option. Adding such an option is important because otherwise my model would predict unrealistically too many wealthy consumers buying used cars.

Now I am ready to introduce my econometric model of the demand system. Let $j \in \mathbf{J} = \{1, 2, \dots, J\}$ be the product index. I assume that a buyer has demand for at most one car. Buyer i 's utility of buying product j at a non-haggling dealership is specified as:

$$u_{ij} = \delta_j (X_{j1}, \xi_j | \theta_1) + \mu_{ij} (X_{j2}, p_j, v | \alpha, \sigma) + \varepsilon_{ij}$$

where

¹³The factor β is assumed to be exogenous to the seller's strategy.

¹⁴The two-step screening mechanism I specified is in essence similar to the simple form of second-degree price discrimination used by Borenstein (1985), except that Borenstein (1985) assumes a reduced form linear cost function for buyers to get the lower price, whereas the disutility of negotiation in my specification is determined endogenously.

$$\delta_j (X_{j1}, \xi_j | \theta_1) = X_{j1} \theta_1 + \xi_j,$$

$$\mu_{ij} (X_{j2}, p_j, v | \alpha, \sigma) = \alpha (\log (y_i - p_j) - \log (y_i)) + \sum_{k \geq 0} \sigma_k v_{ik} X_{j2,k}, \quad v_{ik} \sim N(0, 1), \quad k = 1, \dots, K_2$$

where δ_j is the mean utility of owning product j ,¹⁵ which depends on product characteristics X_{j1} and ξ_j . ξ_j is product j 's quality that is observed by buyers and sellers but not by the econometrician. μ_{ij} is buyer i 's idiosyncratic value for product j . In the above specification, μ_{ij} depends on price p_j , some other product characteristics X_{j2} , buyer i 's income y_i , and the individual specific marginal value $\sigma_k v_{ik}$ for $X_{j2,k}$. $(\theta_1, (\sigma_k)_k, \alpha)$ are the coefficients to be estimated.

Buyer i 's utility of choosing not to buy a car is assumed to be:

$$u_{i0} = \varepsilon_{i0}$$

So, excluding the logit error term, $(\delta_j + \mu_{ij})$ in u_{ij} measures the buyer's net surplus of buying product j relative to the outside option.

By my assumption of the negotiation procedure and the form of haggling dealers' screening strategies, I specify the utility of buyer i buying a product j from a haggling dealership as:

$$u_{ij} = \max \{ \delta_j + \mu_{ij}, \beta (\delta_j + \mu_{ij} + \alpha [\log (y_i - p_j + d_j) - \log (y_i - p_j)]) \} + \varepsilon_{ij}$$

Recall $\{\beta, d_j\}$ defines the dealership-specific screening strategy to be recovered from the data. I point out that in this specification the part of utility measured by the econometric error term ε_{ij} does not affect the buyer's reaction to dealers' sequential offers. I choose such a specification for two reasons. First, if I allow the buyer to incorporate ε_{ij} into her haggling decision I cannot keep the analytical form of the logit individual demand functions. Second, if I interpret ε_{ij} as the transportation cost (or any other cost that is sunk before haggling) from individual i to product j , then it arguably should not enter buyer i 's haggling decision. Theoretical models of competitive price discrimination based on discrete-choice models, e.g. Armstrong and Vickers (2001), take the same approach for tractability.

Let $critical_{ij} = \frac{\beta \alpha [\log (y_i - p_j + d_j) - \log (y_i - p_j)]}{1 - \beta}$, which is the threshold value of net surplus that divides buyers' decisions to accept or reject the seller's initial offer. Then I can rewrite the above utility function for haggling products as:

$$u_{ij} = \begin{cases} \delta_j + \mu_{ij} + \varepsilon_{ij}, & \text{if } \delta_j + \mu_{ij} \geq critical_{ij} \\ \beta (\delta_j + \mu_{ij} + \alpha [\log (y_i - p_j + d_j) - \log (y_i - p_j)]) + \varepsilon_{ij}, & \text{if } \delta_j + \mu_{ij} < critical_{ij} \end{cases}$$

It is clear from the rewritten utility function that in the haggling a buyer accepts the list price if and only if her reservation value $\delta_j + \mu_{ij}$ exceeds the threshold value $critical_{ij}$. Discounts will thus

¹⁵I will also refer to δ_j as the mean quality of product j later.

only be obtained by a selected group of consumers.

Letting $slype_j$ be the indicator function for when the seller of product j is a non-haggling dealership, I can write the entire utility function as:

$$u_{ij} = slype_j \cdot (\delta_j + \mu_{ij}) + (1 - slype_j) \cdot \max \left\{ \delta_j + \mu_{ij}, \beta (\delta_j + \mu_{ij} + \alpha [\log (y_i - p_j + d_j) - \log (y_i - p_j)]) \right\} + \varepsilon_{ij}$$

The heterogeneity in buyer i 's utility comes from $\omega_{ij} = ((v_{ik})_k, y_i, \varepsilon_{ij})$. I assume that $((v_{ik})_k, y_i, \varepsilon_{ij})$ are mutually independent. The set of parameters I need to estimate is $\theta = (\theta_1, \theta_2, \theta_3)$, where $\theta_2 \equiv ((\sigma_k)_k, \alpha)$, $\theta_3 \equiv (\beta, d_j)_j$ ¹⁶.

The market share functions then can be derived given the specification of the utility functions. The set of consumers buying product j is:

$$\Omega_j(p_j, p_{-j} | \theta) = \left\{ \omega_{ij} | u_{ij} \geq u_{ij'}, \text{ for all } j' \right\}$$

I assume that ε_{ij} is drawn from the Type I Extreme Value distribution and ε_{ij} is i.i.d across i and j . So conditional on (v_i, y_i) , the probability of buyer i buying product j has the logit form:

$$P_{ij} = \frac{\exp(u_{ij} - \varepsilon_{ij})}{1 + \sum_{j'} \exp(u_{ij'} - \varepsilon_{ij'})}$$

Finally, the market share of product j can be computed as: $S_j(X|\theta) = \int_{\Omega_j} dF(\omega_{ij}) = \int_{\Omega_j} P_{ij} dF(v_i, y_i)$, where $F(\omega_{ij})$ and $F(v_i, y_i)$ are respectively the probability measures of ω_{ij} and (v_i, y_i) .

3.2 Motivations for the Model of Negotiation

In this subsection I discuss in more detail the motivation for using the above specification of the negotiation process. I based my choice of the above two-offer screening mechanism on two criteria. First, the model should capture the most important mechanism of how the transaction prices are reached starting from the list price. Second, it should be simple enough so that, given the available data, the essential elements in the mechanism can be identified.

With the dealer using the assumed two-offer screening strategy, consumers would self-select into either accepting the initial price or waiting for the discount. Consumers will get the discount if and only if their valuation of the discount exceeds the disutility of negotiation. The price discrimination is completely based on buyers' reservation values. This mechanism of determining the transaction price is consistent with the available empirical evidence about price discrimination in the auto retail market. For example, Goldberg (1996), using Consumer Expenditure Survey data, finds in the new car retail market that the transaction prices are not statistically significantly correlated with any observed demographic characteristics, and that price discrimination seems mostly based

¹⁶The coefficients are divided into three groups for the convenience of the estimation algorithm.

on the distribution of consumers' reservation prices.

In addition, the dealerships' ability to extract surplus out of the transactions is not necessarily weakened by the assumption of their strategies consisting of only two offers. In fact for the special case of monopoly pricing, if the time allowed for negotiation is fixed, applying a two-offer screening strategy actually strengthens the monopolist's ability to extract surplus out of the transaction if it is about to make at least two offers in any case. Coase (1972) conjectured that as the intervals between the offers made by a monopolist shrink to zero, the transaction price reached in negotiation would converge to the monopolist's cost. The validity of this conjecture has been established by Gul, Sonnenschein and Wilson (1986), and Gul and Sonnenschein (1988). Given the time allowed for negotiation fixed, assuming dealers extending only two offers maximizes the interval between offers. The local market that I am focusing on has five sellers. Thus assuming that the haggling dealerships are only making two offers is not necessarily constraining their ability to maximize their profits. In the end, how well the assumed two-offer screening strategy approximates the dealerships' actual strategies depends on how many offers the haggling dealerships normally make.¹⁷

The limited information available in my data determines that I need a very simplified model of the negotiation process. My data only provide information about the sales at each dealership. No information about the actual negotiations is available. So, given the data, inference about the details of the negotiation process is very difficult. My assumption of the negotiation procedure is an obvious simplification of the reality, but the model is identified and it allows me to estimate important factors like the discounts offered by haggling dealerships in the negotiation.

3.3 Alternative Specifications of the Utility Function

For the specification of the utility function, Berry and Pakes (2007) recently developed a random-coefficient discrete-choice model that does not include the logit econometric error term. The model is called a pure-characteristics (PC) demand model as consumers' utility function depends only on product characteristics (observed and unobserved). In principle, I could also base my specification on this model. Two potential benefits could be obtained if I adopted this model. One is that a consumer's decision in the negotiation process would be determined by his/her total utility, instead of the total utility excluding the logit error term as in the above model. The other is that the logit error term, often interpreted as preference for variety, is not included.¹⁸ In my case, it is possible that with the logit error term the relative difference between the same models sold by different dealerships are exaggerated compared to the difference between the different models sold by the same dealership. Thus if sufficient information about the products is observed, the pure-characteristics demand model might be a better choice. An estimation algorithm similar to the one

¹⁷In the simple experiment I personally conducted, the dealer presented me with only one additional offer after telling me the sticker price. The estimated second offer of a product should to some extent capture the average final offers if there are actually offers made between the first and final offer.

¹⁸Whether the exclusion of the logit error term is an advantage or not is an empirical question, as the comparison depends on whether the consumer idiosyncratic preference for the products and diversity is realistic for the specific market in question.

used in Berry, Levinsohn, Pakes (1995) is suggested for estimating the PC demand model. But I find that for my application, the algorithm works much less efficiently for the extension based on the PC demand model than for that based on the BLP model. This particularly is because for the PC model the contraction mapping used in the algorithm to recover the unobserved product characteristics can only be proved to be a weak contraction mapping¹⁹, whereas, the contraction mapping used in the similar algorithm by BLP is strict and has a norm no larger than the total model-predicted market share. I decided not to apply the PC model for my application due to the computational burden of using it.

The problem of using list prices as transaction prices can be partially alleviated by adding dummies for haggling dealers to the utility functions. However the interpretation of this approach is that *every* buyer going to the haggling dealerships gets the discounts, which clearly is not the case. And this specification does not easily accommodate dealers' choices of pricing policies.

4 Data

The data I use for estimation combine several data sets: a dealership-level aggregate retail data set I collected from a car-trading platform website, empirical distribution of household income from the Current Population Survey (CPS), and used car purchase probability by income levels from Consumer Expenditure Survey (CEX). In the following I introduce the dealership-level aggregate retail data in more detail.

Through an online trading platform, Cars.com, I collected the retail data, from August 2008 to August 2009, of list prices, total sales, and the characteristics of all the used cars sold by five dealerships in the White Marsh community in Baltimore county in Maryland. The website, Cars.com, is one of the two major platforms where car dealerships advertise their inventory and buyers get the essential information about the cars before they visit the dealership. Dealers put up all the detailed information about their cars on the website, which includes the information of mileage, engine, transmission, body type, color, features and equipments. Very importantly, all the dealerships advertise their entire inventory on the website.²⁰ The website charges about \$30 per listing per month for individuals and an amount less than that for dealerships depending on the listing volume. Dealerships remove the listings once the cars are sold. I scrape and record the information posted on the listing webpages on a daily basis to get the list prices and car characteristics. I subtract the cars listed on the last day of a period from the pool of all the cars ever listed during the period to get the data of cars sold by each dealership during the period. I get the sales information based on these data of the cars sold. The dealers may decrease or occasionally

¹⁹The norm of the contraction mapping is only shown to be bounded above by 1 in the weak contraction mapping. So, theoretically, the amount of time needed for the weak contraction mapping to find the fixed point does not have an upper bound. This is confirmed by the experience of the simulations, kindly shared with me by Ariel Pakes, conducted in Berry and Pakes (2007).

²⁰CarMax explicitly says in its 2007 annual report that it "lists every retail used vehicle on both Autotrader.com and Cars.com".

increase their list prices.²¹ I take the last list prices observed as the list prices for my estimation.

One important question about my sales data is how accurately my data measures the actual sales by the dealerships. One concern is that some of the listed used cars may be sold through other sales channel, wholesale auctions for example, at prices lower than the list prices. This is most likely to happen if cars stay on the market for too long. The public information I obtained about the dealerships suggests that the problem is very small even though it indeed exists.²² Another concern is that the cars listed at a local store by a chain dealership might be transferred to and sold at other stores of the chain.²³ It may cause problems as I have no information about the location where the listed cars were sold²⁴. Such transfers are mostly done to make up for short-term shortfall in inventories at individual stores. And I expect that the net transfers out of/into the Carmax store in my data to be small as the other Carmax stores in the region are of similar sizes. As I aggregate my data by quarters, the problem caused by such transfers should be small if only similar cars are randomly transferred back and forth between the stores and the net transfers out of/into the local stores in my data are small. Furthermore, the estimated price elasticities I obtained based on the data are very similar to those estimated in other empirical studies of the auto market that did not suffer from the transfer problem. Thus overall I believe that my data do provide good measures of the actual sales by each dealership.²⁵

I also collected some dealership characteristics from the dealership's own websites, which includes pricing policies, quality certification policies, warranty, etc. In my data set, I have two, CarMax and Castle, out of five dealerships formally claim that they do not haggle with buyers. Figure 1.1 and 1.2 are the snapshots of the websites of the two non-haggling dealerships. Both dealerships touted the no-haggle pricing policy as one of their major advantages.

I get the empirical household income distribution from the most recent CPS, and the probability of buying a used car for each of the three income groups from the CEX. Ideally, I need the purchase probabilities for the population of potential buyers during the same period of my data, but the CEX only provides household residence information up to state and normally release the survey data for public use two years after the survey. As a compromise, I use the Maryland state-level purchase probabilities by income levels from CEX 2006, which will be matched with the model-predicted purchase probabilities in my estimation. Accordingly, I use the Maryland 2006 state-level empirical income distribution from CPS.

The yearlong retail data are divided into four quarters in the estimation. Summary statistics

²¹I abstract away from the dynamic aspects of pricing in this paper.

²²Quote from Carmax's 2007 annual report: "Because of the pricing discipline afforded by the inventory management and pricing system, more than 99% of the entire used car inventory offered at retail is sold at retail."

²³This is a potential problem only for one dealer, Carmax, in my data. Other dealerships either only operate one store or don't offer transfers between their stores at the request of buyers.

²⁴Actually, the car history database of CARFAX does have the information about where a car was sold. But I did not have enough resource to get such information for all the cars in my data.

²⁵One might think that my data collecting strategy applies equally well to the new car market, and so I can study my questions in the new car retail market. But the fact is that many consumers place orders in advance at the new car dealerships so that they get the optional features that they want. As I don't observe this part of the sales for the new car retail market, it's not possible for me to use the same strategy to collect good sales data for individual new car dealerships.

of the variables in my data set are presented in Table 1. Table 2 shows the basic market structure in this local market. Essentially, Carmax is the dominant player in the market. The dealers differentiate themselves by focusing on different brands. For example, Al Packer and Schaefer have their focus on American brands, while Carmax and National focus more on Asian brands. I will come back to these two tables when I analyze the possible causes of dealerships choosing different pricing policies.

5 Identification and the Estimation Procedure

5.1 Identification

In this subsection I show the identification of the two-offer screening mechanism specified in my empirical model. Specifically, I show for my empirical model: a) fixing the value of β , the discounts d_j are identified; b) given my data, the joint identification of (β, d_j) is impossible; c) for a larger given β , I would get smaller estimates of the discounts d_j . Furthermore I show that compared to the empirical model I specified, the baseline model necessarily underestimates the coefficient, α , of the price-income interaction term, $\log\left(\frac{y_i - p_j}{y_i}\right)$.

First, when fixing the value of β , the discounts d_j are identified. Figure 2.1 illustrates the idea underlying the identification of d_j when the true β is known. For easier notation, I let $\eta_{ij} = \mu_{ij} - \alpha \log\left(\frac{y_i - p_j}{y_i}\right)$, i.e. η_{ij} is the idiosyncratic part of consumer i 's utility for product j excluding the price-income interaction term. Thus with the additional notation, consumer i 's utility for product j can be conveniently written as

$$u_{ij} = \delta_j + \eta_{ij} + \alpha \log\left(\frac{y_i - p_j}{y_i}\right) + \varepsilon_{ij}$$

if product j 's seller uses posted prices, or otherwise

$$u_{ij} = \max\left\{\beta\left(\delta_j + \eta_{ij} + \alpha \log\left(\frac{y_i - p_j + d_j}{y_i}\right)\right), \delta_j + \eta_{ij} + \alpha \log\left(\frac{y_i - p_j}{y_i}\right)\right\} + \varepsilon_{ij}$$

I show in the following that, if there are sufficiently many no-haggle products in the market, the discounts d_j are identified by matching model-predicted market shares with the observed market shares. The model-predicted market share of product j can be written as²⁶

$$s_j = \int \int \int_{\Omega(\eta_{ij}|y_i, \varepsilon_{ij}, \varepsilon_{ik}, \eta_{ik})} S_j(\eta_{ij}|y_i, \varepsilon_{ij}, \varepsilon_{ik}, \eta_{ik}) dF(\eta_{ij}|\eta_{ik}) dF(\eta_{ik}) dF(y_i, \varepsilon_{ij}, \varepsilon_{ik})$$

where²⁷

$$\Omega(\eta_{ij}|y_i, \varepsilon_{ij}, \varepsilon_{ik}, \eta_{ik}) = \{\eta_{ij}|u_{ij}(\eta_{ij}|y_i, \varepsilon_{ij}) \geq u_{ik}(\eta_{ik}|y_i, \varepsilon_{ik})\}$$

and product k is consumer i 's most preferred choice among all products excluding product j . I

²⁶Recall that (η_{ij}, η_{ik}) is assumed to be independent of $(y_i, \varepsilon_{ij}, \varepsilon_{ik})$

²⁷ $\Omega(\eta_{ij}|y_i, \varepsilon_{ij}, \varepsilon_{ik}, \eta_{ik})$ is just the set of values of η_{ij} that leads consumer i to choose product j .

assume that product k is sold by a non-haggling dealer, so that u_{ik} is fixed given $(\theta_1, (\sigma_k)_k, \alpha)$. It is easy to check mathematically and Figure 2.1 shows that the measure of the set $\Omega(\eta_{ij}|y_i, \varepsilon_{ij}, \varepsilon_{ik}, \eta_{ik})$ strictly increases with d_j . So a larger discount d_j implies a larger model-predicted market share for product j . Therefore there exists a unique d_j that matches the model-predicted market share with the observed market share.

The identification can be seen more transparently in Figure 2.1. Note again that $u_{ik} \geq u_{ij'}$ for $\forall j' \neq j$. The vertical axis measures u_{ij} . The horizontal axis measures $\delta_j + \eta_{ij}$. The horizontal dashed line represents the level of u_{ik} . The heavy solid line is the graph of u_{ij} . The set of consumers that buy product j is: $\Omega(\eta_{ij}|y_i, \varepsilon_{ij}, \varepsilon_{ik}, \eta_{ik}) = \{\eta_{ij} | \eta_{ij} \geq \eta_{ij}^*\}$, where η_{ij}^* is a function of $(d_j, y_i, \varepsilon_{ij}, \varepsilon_{ik}, \eta_{ik})$. $v(d_j | \beta, \alpha, y_i, p_j)$ is the net value of d_j to a consumer i with $\delta_j + \eta_{ij} + \alpha \log\left(\frac{y_i - p_j}{y_i}\right) = 0$, and it is an increasing function of d_j . The figure shows that a larger d_j leads to a larger $v(d_j | \beta, \alpha, y_i, p_j)$ and smaller η_{ij}^* , and thus a larger set of $\Omega(\eta_{ij}|y_i, \varepsilon_{ij}, \varepsilon_{ik}, \eta_{ik})$. If the measure of the kind of cases described in Figure 2.1 is strictly positive, larger d_j then implies strictly larger model-predicted market share of product j , and there exists a unique d_j that satisfies the market share equations. Thus if there are sufficiently many no-haggle products, the discount d_j is identified given that β is known. As there are many products sold under posted price in my data, the condition can be met without too much difficulty. It is worth noting that if no dealer offers posted price, only the differences between the discounts offered by the dealerships can be identified.

Second, given my data, the joint identification of (β, d_j) is impossible. The idea of showing this claim is presented in Figure 2.2. The figure shows the utility function u_{ij} under two sets of parameters, (β, d_j) and (β', d'_j) . Both of them result in the same value of η_{ij}^* , and thus the same model-predicted market share for product j . Essentially, I cannot identify both the slope and intercept of a line with the knowledge of only one point on the line. Thus, given that I only have the information about dealers' aggregate sales, joint identification of (β, d_j) is impossible. But I do know how the identified d_j would change with the specified value of β . The lighter solid line gives the utility function for some $\beta' < \beta$. To fit the η_{ij}^* corresponding to the observed market share, it is necessary that $d'_j > d_j$. Thus, the estimated d_j would be smaller if a larger β is specified. Intuitively, a smaller β' implies larger costs of negotiation to buyers. When the discount is harder to get, there must be a larger discount obtainable in order to attract the same number of buyers.

Third, the baseline model necessarily underestimates the coefficient, α , of the interaction term $\log\left(\frac{y_i - p_j}{y_i}\right)$. The idea of the argument for the claim is shown in Figure 2.3. Controlling for the other factors, the prices listed by haggling dealers are higher than those listed by the non-haggling dealers. As the baseline model treat all list prices as the transaction prices, the consumer's disutility for prices in the model would have to be lowered in order to fit the observed haggling dealers' market shares. This can be seen clearly in figure 2.3. Given the true value of α , the baseline model would predict that consumers in the set of $\{\eta_{ij} > \bar{\eta}_{ij}\}$ would buy product j , which is smaller than $\{\eta_{ij} > \eta_{ij}^*\}$ implied by the data. Forcing the baseline model to fit the observed market shares would drive down α . This is because, as α becomes smaller, $\delta_j + \eta_{ij} + \varepsilon_{ij} + \alpha \log\left(\frac{y_i - p_j}{y_i}\right)$ increases

faster than $\delta_k + \eta_{ik} + \varepsilon_{ik} + \alpha \log\left(\frac{y_i - p_k}{y_i}\right)$ given that $p_j > p_k$, and thus it moves $\bar{\eta}_{ij}$ towards η_{ij}^* . In the figure, fitting the baseline model to the market share data would drive α down to α' .

5.2 Estimation Procedure

Following Berry, Levinsohn, Pakes (1995) and Petrin (2002), I estimate my model by using the General Method of Moments (GMM). I use three sets of moments to identify the structural parameters. The first set of moment conditions are the equations of market shares that match the model-predicted market shares to those observed in my data.

$$s_j = S_j(\delta_j, \delta_{-j} | \theta_2, \theta_3)$$

I get the mean product qualities δ as a function $\delta(s, \theta_2, \theta_3)$ of (s, θ_2, θ_3) by inverting the above market share functions. Because the inverse functions are not analytical, I use the contraction mapping suggested by BLP to compute the inversion. Although my specification is not the same as that in BLP, it is easy to show that the mapping used in BLP is still a contraction mapping under my model due to the fact that $\beta < 1$, and the norm of the contraction mapping is at most “1—the model-predicted share of the outside option”²⁸.

Second, as is now well-known, the product prices could be correlated with the structural error term ξ_j because of the simultaneous determination of product prices and market shares: the unobserved quality ξ_j could have been taken into consideration when sellers set prices. I follow BLP in choosing my instrumental variables IV_j for price, assuming that $E(\xi_j | IV_j) = 0$. The orthogonal conditions are my second set of moment conditions. The instruments I use are the product characteristics excluding price, the sum of these product characteristics of cars sold by the same dealership, and the sum of these product characteristics of cars sold by the other dealerships. Since the price competition in the used car retail market is among dealerships, I construct my instruments by summing over dealerships. The order condition of identification is that the number of instruments must be no less than the number of endogenous variables²⁹. I have more than enough instruments for my estimation.

To get more information for estimating the parameters in μ_{ij} , I use a third set of moments that matches the model-predicted probabilities of purchasing a used car by income levels to those computed using CEX data. As it is impossible to simultaneously identify β , I set $\beta = 0.98$ ³⁰. $(\delta_0, \theta_1, \theta_2, (d_j)_j)$ are the coefficients that I am going to estimate.

Combining the above three sets of moment conditions, I estimate the parameters θ by minimizing the following objective function:

$$f(\theta) = \left(\begin{array}{c} IV' \cdot (\delta(\theta_2, \theta_3) - X_1 \theta_1) \\ P(\theta | X) - prob \end{array} \right)' W \left(\begin{array}{c} IV' \cdot (\delta(\theta_2, \theta_3) - X_1 \theta_1) \\ P(\theta | X) - prob \end{array} \right)$$

²⁸The estimation algorithm can run very slow if the initial values used by algorithm predicts very small share for the outside option.

²⁹That is “one plus the number of parameters in (θ_2, θ_3) ”.

³⁰Later I will discuss the robustness of the estimation results to the variation in the value of β .

where $P(\theta|X)$ is the vector of model-predicted purchase probabilities by income levels, and $prob$ is the vector of purchase probabilities obtained from CEX, and W is the optimal weighting matrix:

$$W = E \left(\begin{pmatrix} IV' \cdot (\delta(\theta_2, \theta_3) - X_1\theta_1) \\ P(\theta|X) - prob \end{pmatrix} \cdot \begin{pmatrix} IV' \cdot (\delta(\theta_2, \theta_3) - X_1\theta_1) \\ P(\theta|X) - prob \end{pmatrix}' \right)$$

I refer readers to Petrin (2001) for the details of computing W .

6 Estimation Results

6.1 The Demand System

I will discuss my estimation results from the following two perspectives. First, whether my augmented model is necessary to estimate the demand system. Specifically, how big a difference the augmented model makes in terms of the estimated coefficients and price elasticities. Second, how do the estimated discounts compare to the actual discounts being offered by the dealerships, and what are the sources and direction of the potential bias in my estimates.

The estimation results are reported in Table 4. I compare my results to those estimated under the baseline model and logit IV model. The baseline model is nested within my augmented model: it is equivalent to restricting all the discounts from haggling dealerships to be zero. In principle, when the posted-price assumption does not hold, the coefficients estimated in the baseline model would be significantly different from those estimated in my augmented model. Let $\theta_{12} = \begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix}$, i.e. the parameters shared by the baseline model and the augmented model. I test the significance of the difference between the two models by constructing the following Wald statistic through bootstrapping³¹:

$$Wald = \left(\hat{\theta}_{12} - \tilde{\theta}_{12} \right)' \left(Var \left(\hat{\theta}_{12} - \tilde{\theta}_{12} \right) \right)^{-1} \left(\hat{\theta}_{12} - \tilde{\theta}_{12} \right)$$

where $\hat{\theta}_{12}$ and $\tilde{\theta}_{12}$ are respectively the coefficients estimated in the baseline model and the augmented model. The Wald statistic has a chi-squared distribution with the degree of freedom equal to the number of parameters in θ_{12} . The test rejects with probability of essentially one the null hypothesis of the two models being effectively the same. This is expected as the coefficients of price estimated under the two models are very different: the assumption of universal posted price leads to significant smaller estimate of the price effect.

The estimated own price elasticities under the augmented model are in the range from around 3 to 7, while those estimated under the baseline model are generally below 3. Figure 3 presents the own price elasticities computed based on the two estimated models. It shows that the assumption of universal posted price leads to significant smaller estimates of own price elasticities, consistent with my analysis in the identification subsection. The intuition for this result is straightforward. The baseline model links the list prices to the observed market shares, which were actually generated

³¹For bootstrapping, I use 200 random samples drawn with replacement from the original data of all the cars sold by the five dealerships. The bootstrap samples are of the same sizes as the original data sample.

by the possibly lower transaction prices. In the estimation, this linkage forces the underestimate of the disutility of price. Similar problems with the observed prices exist in many other markets as well. For example, car manufacturers offer large discounts late in each fiscal year through clearance sales, but the information about the discounts is often missing and the prices observed by the econometrician are often only the manufacturer suggested retail prices or some wholesale prices. In the grocery product market, the observed prices (e.g. shelf prices) are also possibly higher than the actual transaction prices as discounts are given to some targeted consumers. My results suggest that direct uses of such observed prices in empirical research would lead to smaller estimates of the price elasticities, and my method represents one possible way to deal with the measurement problem in the observed prices.

The point estimates of the discounts for the three haggling dealerships are respectively \$1969, \$733, and \$2080 in the model with $\beta = 0.98$. The 95 percent confidence intervals for the three discounts are respectively (\$844,\$3094), (-\$137,\$1603) and (\$1217, \$2942). These estimates are in general consistent with anecdotal evidence. But to get a better sense of how the estimated discounts compare with the actual discounts being offered by dealerships, I compare my estimates to the discounts obtained in previous field experiments. Table 6 presents the computed discounts obtained by the testers in the experiment³² referenced in Ayres and Siegelman (1995). The experiment was carried out in new car dealerships in the Chicago area in 1990. The average cost of a new car in 1990 was about \$16,000, which is close to the average list price in my data. The weighted average discounts obtained by the testers is \$1184, which is within two of the 95 percent confidence intervals above. The average discount obtained by the majority male testers is \$1476³³, within all three 95 confidence intervals above. Thus the magnitude of the estimated discounts seem to make sense. To get further direct evidence about the actual discounts offered by the dealers in my data, I and two friends, posing as used car buyers, went to the dealerships in my data. I seriously negotiated for a few deals. For example, I looked at a 2005 Chrysler Pt Cruiser with the list price of \$10,995. I went through the whole process that an actual potential buyer would have gone through, including test driving. The negotiation results are presented in Figure 4. The tentative offer I obtained for the car was \$9995. The \$1000 discount is within the 95 percent confidence interval, (\$844,\$3094), estimated for the dealer. Thus, overall the evidence I can find about the actual discounts given by dealerships is consistent with my estimates. The discount estimates and the corresponding dealer profits are part of the key information that I use later in my comparison of pricing strategies.

In Table 5 I compare the coefficients estimated with various values of β , and the coefficients estimated by simply adding dummies for the haggling dealerships to the baseline model. It is easy to see that as β converges to 1, the augmented model converges to the baseline model with dummies for haggling dealers. The estimation results show that the estimated discounts decrease as the specified value of β becomes larger, which is consistent with what I have shown about the relation between the specified β and the estimated discounts. The analysis I do later will be based

³²For the details of the experiment, see Ayres (1991) and Ayres and Siegelman (1995).

³³As I used initial offer prices for black male testers as list prices, the actual discounts off the list prices can be a little larger than those provided in Table 6.

on $\beta = 0.98$. Without the knowledge of the true value of β , misspecified β is a source of bias in my estimates of discounts.

6.2 Further Discussions

Other factors affecting the transaction prices that my model does not fully capture could also lead to bias in my estimates. For example, Scott-Morton, Zettelmeyer, and Silva-Risso(2004) find that car buyers who are patient³⁴ paid 0.47 percent less than others. Yet I let the discounting factor for all consumers be fixed at a given value of β . The best I can hope for is that the β I focus on captures average patience.

Another important aspect that may have been missed by my model is the impact of the information actively collected by consumers on transaction prices. For example, two such factors affecting transaction prices tested by Scott-Morton et al (2004)³⁵ are buyers' knowledge of dealers' invoice prices and their own outside options. The first measures buyers' information of dealers' reservation values, while the second measures consumers' threat points in negotiation obtained through searching. Their findings are that better consumer information does decrease the transaction price. But, the combined total effect of the two information factors equals to about one half of the total effect of consumers' disutility for bargaining, and statistically the impacts of both information factors are less significant. My model does not seem able to correctly capture the impacts of information on transaction prices, and this may lead my model to overestimate the actual discounts obtained by consumers. For example, it is possible that consumers with better information could obtain some discounts incurring only negligible disutility.³⁶ One way the buyer may obtain an instant discount is, for example, by presenting some credible evidence (such competing offers) about an upper bound of her reservation value, since the seller may want to lower its initial offer given the updated upper bound of the distribution of the buyer's reservation value. As I have shown in the identification subsection, if the specified value of β is smaller than the true value of β , I would overestimate the magnitude of the discounts³⁷. Thus if part of the discounts were obtained by consumers without incurring disutility, i.e. as if β is actually 1, then my model would overestimate the discounts. I will come back to this point when I discuss the results of comparing dealerships' profits under various pricing policies. Admittedly, without the information about the exact channels that are not captured by my specification, it is hard to more rigorously discuss the potential bias in my

³⁴The buyer patience in their paper was measured by whether consumers choose the answer "I would not have bought a car at that time" when asked about their alternative course of action had negotiations with the dealer broken down.

³⁵Their study uses data from California new car market. The market, the location and thus the consumer population could be quite different from what our data represents. The difference should be kept in mind while consulting their estimated impacts of various factors on transaction prices.

³⁶This is supported by the experience of a recent experiment conducted by us personally. When we told the dealer that we had a few competing offers in hand, the dealer quickly presented their offer with a large discount off the list price, yet in another incidence at another dealership, when we presented no such information, the dealer did not budge.

³⁷For example, when β changes from 0.98 to 0.95, the estimated discount of the first dealership (Al Packer) increases by about \$530, and that of the fifth dealership (Schaefer) increases by around \$520.

estimates of discounts.

One might also be concerned about the possibility that the estimated discounts may in part reflect other unobserved heterogeneity across dealerships. While this is in principle possible, it is unlikely an important factor affecting my estimates. The dealers in my data set are located within 3 miles of each other, and before I collected the data, I tried to make sure that aside from the cars sold by the dealerships, they are otherwise similar. In addition I deducted the market values of dealer specific effects from the list prices whenever possible. For example, cars sold by some dealers come with a warranty of 3 month, or 3000 miles. The market value of such extra features are small, and I deducted the market values of such features from the list prices to make cars comparable across dealerships. In addition, in my estimation, I used instruments to take care of the potential product-level unobserved heterogeneity. So, I recognize that there are potentially some unobserved dealership-specific effects, but I believe that the values of such effects are small and that the estimates of the discounts mostly reflect the haggling dealers' unobserved discounts, instead of other dealership specific effects.

One other concern about the model I estimated is that it may be restrictive for me to assume identical discounts for all the cars sold by a dealership. One way to check the empirical relevance of this concern is to divide the models sold by each dealership into several groups, estimate a discount for each group, and check if the estimated discounts vary significantly across the groups. Defining such groups by list prices seems reasonable for my purpose. More specifically, I let each dealership give two possibly different discounts (d_1, d_2) : one for low-end cars, one for high-end cars³⁸. I test whether the discounts being offered actually vary significantly with prices by testing the null hypothesis of $d_1 = d_2$ for each dealership using the following test statistics:

$$W = \left(\hat{d}_1 - \hat{d}_2 \right) \left(\text{Var} \left(\hat{d}_1 - \hat{d}_2 \right) \right)^{-1} \left(\hat{d}_1 - \hat{d}_2 \right)$$

The test results do not reject the null hypothesis for any of the three haggling dealerships. This suggest that requiring the discounts to be the same for all cars sold by the same dealership is not that restrictive empirically. Furthermore, in the counterfactual experiments to be introduced later, the counterfactual optimal discounts computed for smaller components of Carmax are very similar.³⁹ It shows for example that even if I allow a large dealership like Carmax to costlessly set different discounts by car quality, it would only have relatively small effect on the total profit. Thus considering the practical costs of screening buyers and the relatively small gains after allowing for more flexible discounts, I believe the empirical results of this paper will remain the same even if I allow dealerships to set more flexible discounts.

³⁸Cars are defined to be low-end cars iff their prices are less than \$13000, which is the average list price.

³⁹The counterfactual experiments will be introduced in detail later.

6.3 Estimating Marginal Costs

In this subsection, I estimate the marginal costs of the products. The marginal costs will be used later in computing dealerships' profits. To evaluate how well my model captures the actual competition in the market, I will compare the estimated price-cost margins with a set of market average price-cost margins I compiled using external information.

To estimate the marginal costs, I assume that the pricing policies are exogenous in the short run⁴⁰ so that the dealerships compete only in list prices, and that my sales data are generated by a Nash Equilibrium in list prices. I also assume that the marginal costs of all products are constant. Given these assumptions, in the following I use the first order conditions in dealers' profit optimization problems to recover the products' marginal costs.

Let $L = \{1, \dots, l, \dots, \bar{l}\}$ be the set of dealerships, J_l be the set of products sold by dealership l , M be the total market size, and ρ_j be the proportion of buyers of product j that get its discount. $\rho_j = 0$ for products sold by non-haggling dealerships. Then dealership l 's profit maximization problem can be written as:

$$\max_{\{p_j\}_{j \in J_l}} M \cdot \sum_{j \in J_l} ((p_j - mc_j - \rho_j \cdot d_j) S_j(p_j, p_{-j}))$$

Note here dealers are not discounting the profits they receive after negotiating with consumers, i.e. the negotiation does not add any costs to dealers. The first order conditions of the problem are:

$$S_j + \sum_{j' \in J_j} (p_{j'} - mc_{j'} - \rho_{j'} d_{j'}) \frac{\partial S_{j'}}{\partial p_j} - \sum_{j' \in J_j} d_{j'} S_{j'} \frac{\partial \rho_{j'}}{\partial p_j} = 0, \text{ for } j \in J_l$$

These conditions can be written concisely in matrix form as the following:

$$S + \Delta_p S \cdot (p - mc) - \Delta_p S \cdot \bar{\rho} d - \Delta_p \rho \cdot \bar{d} s = 0$$

where $\Delta_p S_{(j,j')} = I_{\{j' \in J_j\}} \cdot \frac{\partial S_{j'}}{\partial p_j}$, $\Delta_p \rho_{(j,j')} = I_{\{j' \in J_j\}} \cdot \frac{\partial \rho_{j'}}{\partial p_j}$, $\bar{\rho} d = (\rho_1 d_1, \dots, \rho_J d_J)'$, $\bar{d} s = (d_1 s_1, \dots, d_J s_J)'$, and $I_{\{j' \in J_j\}}$ is the indicator function for the logical evaluation of $\{j' \in J_j\}$. So, I have the following formula for recovering the marginal costs:

$$mc = p - \bar{\rho} d + (\Delta_p S)^{-1} (S - \Delta_p \rho \cdot \bar{d} s)$$

where $\bar{\rho} d - (\Delta_p S)^{-1} (S - \Delta_p \rho \cdot \bar{d} s)$ is the markup that can be computed by using the estimated demand system. In my static model, the marginal costs I estimate represent the current acquisition costs for each dealership. In a simple dynamic model without uncertainty, the estimated mc should be understood as the current opportunity cost, which would depend on the acquisition cost in the next period and the interest rate faced by the dealer. In the following, I will abstract away from dynamic considerations and assume that mc is just the static unit acquisition cost.

⁴⁰This is an appropriate assumption because switching pricing policies entails nontrivial changes and lots of trials.

To evaluate how well the estimated model captures the competition in the market and identify potential weaknesses in my model, I compare the estimated price-cost margins to another set of measures of the price-cost margins I constructed by using external information. This other set of price-cost margins are constructed based on the information provided by Kelley Blue Book (hereafter abbreviated as KBB). I get from KBB⁴¹ the “suggested retail prices” and “trade-in values” for all the models in my data. These statistics are compiled by KBB based on their survey of car dealerships. The former represents the “representative of dealers’ asking prices and is the starting point for negotiation between a consumer and a dealer”, and the latter is “what consumers can expect to receive from a dealer for a trade-in vehicle assuming an accurate appraisal of condition”⁴². So the two measures matches respectively the list prices and the dealers’ marginal costs in my model very well. Using the external information, I calculate a new set of price-cost margins (I call them the KBB margins hereafter) for each car model in my data using the following formula:

$$\text{KBB Margin} = \text{suggested retail prices} - \text{trade-in values} - \text{average discount}$$

To calculate the KBB margins, I use the “suggested retail price” and the “trade-in values” for models in excellent condition, with the average age, mileage⁴³ and standard options, and I set the market average discount to be 500 dollars⁴⁴. I then compute the market-share-weighted average KBB price-cost margin for each dealership, and compare them to those estimated in my model. As the KBB survey covers dealerships of various sizes, the computed KBB average margins for the dealerships should be taken as the average margin obtainable by dealerships with the average market power selling the portfolio of models carried respectively by each dealership. And I should expect dealerships of larger size get higher margins due to their relatively larger market power. The four dealerships other than Carmax are of around the market average size, whereas Carmax is the largest used car dealership in the US. So I should expect that the four dealerships obtain about market average margins and Carmax obtains higher margins than the market average ones. The two sets of dealerships’ weighted average margins are presented in Table 6. For the three haggling dealerships, the absolute difference between the estimated average margin and the KBB average margin is less than \$100. The estimated average margin for Castle, one of the non-haggling dealerships, is smaller than the corresponding KBB average margin by about 400 dollars. The estimated average margin of Carmax is much larger than the corresponding KBB average margin.

Thus overall my augmented model did very well in predicting the dealership average margins for the four dealerships other than Carmax. The augmented model passes the basic test by correctly predicting that Carmax’s average margin is larger than that of an average dealership carrying the

⁴¹Kelley Blue Book is a very well-known information supplier that provides market information about new and used cars.

⁴²The quotes are from KBB’s website, www.kbb.com.

⁴³In my data, the average car age is about 2.9 years and the average mileage is about 41800 miles.

⁴⁴The 500 dollars seems to me a good guess of the average discount offered in the used car market. Note that it is a average across both uniform pricing and bargaining dealers.

same portfolio of models. But unfortunately, I do not have the necessary information to gauge how much higher the average margin for a dealership like Carmax should be. I suspect that my model may have overestimated the average margin for Carmax for the following two reasons. First, I may have excluded from my data some of the Carmax store’s competitors. My data covers dealerships located within 3 miles from the Carmax store⁴⁵. Yet the Carmax store may cover an area larger than what is covered by my data, and thus some of Carmax’s less direct competitors are not included in my data⁴⁶. So, my model could have overestimated Carmax’s average margin by inflating its market share and market power. Second, the market power enjoyed by Carmax by carrying a very large portfolio of different models might be measured with some error by my model. Yet this source of bias could be much less important as my model predicted the average margin for the other four dealerships with very good precision.

Although the average margin of Carmax may have been overestimated somewhat, the overall performance of the model is decent. So, I can be relatively assured about the validity of my model and proceed to evaluate the performance of alternative pricing policies for each dealership.

7 Counterfactual Experiments

7.1 The Comparison of Profits under Posted Price and Hagglng

Now with the recovered marginal costs, the estimated demand system and the discounts, I am ready to compare the profits of each dealership under various pricing policies. Specifically, I first calculate the estimate of each dealership’s actual total profit. Then, given the pricing policies chosen by the other dealerships, I compute for each dealership the profit under the counterfactual optimal screening strategy.⁴⁷ The computed optimal screening strategies represent the pricing policies that the dealerships should have incentives to switch to if they are not already using them. I also compute the optimal posted prices, restricting the discounts to be zero, for the hagglng dealerships. I compare the computed counterfactual optimal profits and pricing policies with the observed (estimated) ones for each dealership to reveal the relative performance of the various pricing policies.

The comparison results for the five dealerships are presented in Table 8.1-8.5. Let us first look at the comparison of profits under optimal posted price and under the optimal screening strategy. For Carmax, the computed counterfactual optimal discount is around \$602, and applying the optimal discount would increase Carmax’s total profit by about 1.01 percent. For Castle, the computed counterfactual optimal discount is around \$760, and applying the optimal discount would increase

⁴⁵Other used car dealerships are mostly located more than 8 miles away from the Carmax store.

⁴⁶My original intention was to collect the data in a more isolated and well-defined market. As I have shown previously, my identification relies on the existence of sufficiently many no-haggle products, so my strategy was to collect data for markets that have been entered by Carmax for more than 3 years (it takes 2 to 5 years for the business of a new Carmax store to mature). The White Marsh market in Baltimore, Maryland, is the most well-defined such market that I could find.

⁴⁷That is the optimal list prices and discount. I am assuming that each dealership uses the same discount in all its screening strategies.

Castle's total profit by about 2.2 percent. These numbers suggest that given the other dealerships' current pricing policies, the performance of posted price for the two non-haggling dealerships is indeed very close to that of the optimal screening strategy. For the three haggling dealerships, Al Packer, National, and Schaefer, switching from optimal posted prices to the optimal screening strategies increase their profits respectively by 4.7 percent, 4.9 percent, and 7.1 percent. These returns of switching from posted price to optimal screening are much larger than those for the two non-haggling dealerships.

Table 8.1 and 8.4-8.5 also present the estimated actual profits and discounts for the three haggling dealerships. For National, the estimated discount, \$733, is very close to the computed optimal discount, \$628, though the estimated profit obtained by National could be increased by around 10 percent if it chose the optimal discount. Al Packer's estimated actual profit is lower than both that under the optimal posted prices as well under the optimal screening strategy. The underperformance is mainly due to the fact that the estimated discount offered by Al Packer is too big: the estimated discounts is about \$800 higher than the computed optimal discount (\$1151). It could have increased its profit by 20.7 percent if it offered the optimal discount. The results of Schaefer show that the estimated discount chosen by Schaefer is unsatisfactory for the same reason.⁴⁸ It could have increased its profit by around 25 percent if it offered the right discount. These comparisons seem to suggest that the actual performance of the three haggling dealerships could be greatly improved, however as my estimates possibly exaggerated the discounts offered by the dealerships⁴⁹, the discounts actually offered by the three haggling dealerships can be closer to the optimal discounts than the estimates.

In short, the above results show that the relative performance of posted price and haggling vary across dealerships: dealerships using posted price would only see their profits increase slightly if they haggled, whereas the haggling dealerships' profits would drop significantly if they switched to posting prices.

7.2 Competition and the Choice of Pricing Policies

The important question remains after the above analysis is why Carmax and Castle chose and indeed should choose posted price, whereas the other dealerships should apply price discrimination. The evidence I will present in the following suggest that the dichotomous choice of pricing policies is mainly a result of the asymmetric market conditions: the two non-haggling dealerships sell larger numbers of models, and Carmax is more vertically differentiated from the other dealerships. I will argue that under the two conditions the two non-haggling dealerships face less competition and thus would not see significant profit increase if they switched from posting price to haggling. In the following, I first present the theoretical results regarding the choice of pricing policies from the monopoly pricing literature. Then I complement it with two pieces of evidence based on

⁴⁸It's worth noting that the consistent estimates of discounts in step 1 for dealer 1 and 5 are respectively 1.04 and 1.43 thousand dollars.

⁴⁹We pointed out this possibility when discussing the estimated discounts.

counterfactual experiments.

First, theoretically, if price discrimination imposes the same costs on the dealerships as it does on consumers, it can be proved in a static monopoly pricing model that posted price is in fact the optimal pricing policy. For example, the following is the mechanism design problem faced by a monopolist trying to use direct mechanism $(p(v), \phi(v))$ to sell a product to a consumer with unobserved reservation value v :

$$\begin{aligned} & \max_{p(\cdot), \phi(\cdot)} \int_{\underline{v}}^{\bar{v}} (p(v) - c) \phi(v) dF(v) \\ \text{s.t. } IC & : (v - p(v)) \phi(v) \geq (v - p(v')) \phi(v') \\ IR & : (v - p(v)) \phi(v) \geq 0 \\ \phi(v) & \in [0, 1], \forall v \in [\underline{v}, \bar{v}] \end{aligned}$$

where $p(v)$ and $\phi(v)$ are respectively the price and probability of selling to a consumer reporting reservation value v ; c is the seller's cost of producing the product; $F(v)$ is the cumulative distribution function of v ; \bar{v} and \underline{v} are respectively the upper and lower bound of v 's support; the consumer's valuation for the outside option is 0, and *IC* and *IR* respectively stand for the Incentive Compatibility condition and the Individual Rationality condition. It can be proved that, under a relatively mild condition on F , the above program has the following form of solution⁵⁰:

$$\begin{aligned} \phi(v) &= \begin{cases} 1, & \text{if } v \geq v^* \\ 0, & \text{if } v < v^* \end{cases} \\ p(v) &= \begin{cases} v^*, & \text{if } v \geq v^* \\ 0, & \text{if } v < v^* \end{cases} \end{aligned}$$

where v^* is a constant to be solved. This means that the optimal pricing policy for the monopolist is to charge the posted price of v^* . Riley and Zeckhauser (1983) show with similar specifications in a dynamic monopoly pricing model that posted price (called take-it-or-leave-it price in their article) is optimal in the dynamic environment as well.⁵¹ These theoretical results suggest that if a seller is close to being a monopolist, the posted price format may also be its best pricing policy approximately.⁵² It is worth noting that discontinuous solutions like the above non-haggling solution are mainly a result of the fact that the objective function and the constraints are linear in the control function ϕ . This class of optimal control problems is often called the ‘‘Bang-Bang control’’ problems. Without the linearity in ϕ , mechanism design problems like the above one normally have more smooth solutions for (ϕ, p) .⁵³ This means that the optimal solution would

⁵⁰The proof is shown in the Appendix.

⁵¹I am not aware of any article in the literature that explicitly shows how oligopolistic competition affects the equilibrium choice of pricing policies.

⁵²If there exist additional fixed costs to implementing price discrimination, the posted price format can be the actual optimal pricing strategy for sellers with large market power.

⁵³This is because the optimal ϕ 's are normally solutions to ordinary differential equations in ϕ .

involve price discrimination without the linearity in ϕ .

For my empirical model with five competing dealerships, Table 9 summarizes the basic market structure and the proportional returns of switching from posted price to haggling for each dealership. It shows that the two non-haggling dealerships supply the largest number of models, have the largest total market shares, and have the *smallest* proportional return if switched from posted price to haggling. Figure 5.1 shows the empirical cumulative distribution of the estimated mean qualities of models supplied by each dealership.⁵⁴ The distribution of the qualities supplied by the four dealerships other than Carmax are very similar to each other, however, the estimated mean qualities supplied by Carmax significantly exceed those supplied by the other dealerships. Generally, we expect a seller to have larger market power when the qualities supplied by the seller are more differentiated from those supplied by the other sellers, or when the seller carries more models relative to the other sellers. The patterns shown in Table 9 and Figure 5.1 seem to suggest that the market power derived in part from the above two sources reduces dealerships' incentives to haggle.

I further investigate the hypothesis through the following two sets of counterfactual experiments. The basic idea of the experiments is to see how the comparison between posting price and haggling would change with the competition faced by a dealership. I employ the above two sources of variation in competition: the vertical differentiation in cars' quality and the number of models carried by a dealership. In the experiments, I create counterfactual dealerships by using the pool of models supplied by Carmax. For the first set of counterfactual experiments, I divide the pool of models carried by Carmax into N groups to form N Baby Carmaxes. The n th Baby Carmax has the mean qualities of all its models in the range from the $(100(n-1)/N)$ th percentile to the $(100n/N)$ th percentile value of the empirical distribution of mean qualities of all models supplied by Carmax. I compute the profits for each Baby Carmax under the optimal posted prices and under the optimal two-offer screening strategy while fixing the posted prices of the rest of the models supplied by Carmax and the list prices and discounts of the models sold by all the other dealerships. The exercise is similar to that done for the original Carmax except that the counterfactual profits here are computed for the Baby Carmaxes.⁵⁵ As now Carmax is decomposed into N Baby Carmaxes, I am able to identify how the components are affecting the overall incentives in the Carmax's choice of pricing policies. In my experiments, I set $N = 12$.⁵⁶

In the second set of experiments, I create counterfactual dealerships carrying various numbers of models. Specifically, I create a counterfactual dealership carrying J models by randomly sampling J models without replacement from the pool of models carried by Carmax. Then I compute the counterfactual profits for the Baby Carmaxes under the optimal posted prices and under the optimal screening strategy in the same situation as described above. I repeat the experiment 50 times for each given J , and compute the average return when the counterfactual dealerships of size J switch

⁵⁴Recall that δ_j specified in the utility function is the mean quality of product j .

⁵⁵One additional difference is that in the experiment, each Baby Carmax also faces the competition from the rest of the models sold by the original Carmax.

⁵⁶I also tried $N = 9$, which give me very similar results.

from posting price to haggling. I do the exercise for $J = 11, 15, 21, 43$, which respectively amount to around $1/12^{th}$, $1/9^{th}$, $1/6^{th}$, and $1/3^{rd}$ of the number of models carried by Carmax.

The results for the first set of experiments are presented in Table 10.1 and plotted in Figure 5.2. It is clear that when Baby Carmaxes switched from posting price to haggling, profits increase most significantly for Baby Carmaxes supplying models with qualities that are most frequently supplied by the other dealerships. In contrast, the Baby Carmaxes carrying models with qualities that are rarely supplied by the other dealerships (that is at the two ends, especially the high end) would only see their profits increase slightly if they switched to haggling from posting price. These results suggest that when a dealership is more vertically differentiated from the other dealerships, haggling becomes a less appealing pricing policy. The results for the second set of experiments are presented in Table 10.2. The results show that when the number of models carried by the Baby Carmaxes become smaller, the return of switching from posting price to haggling on average becomes larger. In my empirical model, a seller selling more models has larger market power, because consumers have preferences for variety as captured by the logit error term in the utility function. Therefore, results of both two sets of experiments are consistent with the hypothesis that more intense competition tends to make haggling more appealing relative to posting price.

I am not aware of any article in the literature that shows directly how in an oligopoly the increase in competition affects the equilibrium choice of pricing policies. My intuition for the effect of competition on the relative advantages of posting price versus haggling is the following. In the monopoly pricing case, supplying to additional buyers with lower reservation values is costly because it would increase the informational rent to all the inframarginal buyers with higher reservation values. The *constant* cost of selling to all buyers makes it optimal to sell only to buyers with reservation values above a threshold with certainty⁵⁷. Yet in the situation with differentiated products and multiple sellers, each seller faces competition from other sellers. As it is potentially more profitable to sell to buyers with higher reservation values, the competition for buyers with higher reservation values would be more intense relative to the competition for buyers with lower reservation values. The competition would thus reduce the payoff of selling to buyers with higher reservation values by a larger amount relative to selling to buyers with lower reservation values.⁵⁸ With competition, when considering whether to sell to additional buyers with lower reservation values, the increase in the informational rent to buyers with higher reservation values becomes a smaller problem because now, due to competition, part of the increased informational rent cannot be retained by the seller in any case. Thus selling to buyers with lower reservation values may become profitable even at the cost of the increased informational rent to buyers with higher reservation values. Therefore, with more competition, haggling may become more appealing relative to posting price.

Technically, the existence of competition in the static environment would make a seller's objective function nonlinear in its control function ϕ . The objective function of a seller may take the

⁵⁷See the appendix for more details.

⁵⁸With regard to a seller's payoff function, this in effect is similar to changing the constant cost of selling to all buyers to increasing cost of selling to buyers with higher reservation values.

following form

$$\max_{p(\cdot), \phi(\cdot)} \int_v^{\bar{v}} (p(v) - c) H(\phi(v)) dF(v)$$

where $H(\phi(v)) < \phi(v)$, and $H(\phi)$ depends on the pricing strategies of the competitors and is nonlinear in ϕ . This change in the objective function would make the optimal solution unlikely to take the above non-haggling form. A rigorous theoretical analysis of why sellers may choose different pricing policies in equilibrium warrants a separate article. I leave it for my future research.

8 Concluding Remarks and Future Work

This paper incorporates actual pricing policies into the popular BLP framework. With the augmented model, I am able to estimate the dealerships' actual profits and the demand system using data with only list prices. My estimates of the unobserved discounts are the same in magnitude as those found in previous experiments as well as in actual buying experiences at these dealerships. Using the estimated demand system, I computed the dealerships' profits under various pricing policies. Comparing these profits I found that dealerships using posted price would only see their profits increase slightly if they haggled, whereas the haggling dealerships' profits would drop significantly if they switched to posting prices. Through two sets of counterfactual experiments, I also found that the return to switching from posting price to haggling increases with the competition faced by the individual dealerships. Combined with results in the monopoly pricing theory, the experimental evidence suggests that competition is actually encouraging dealers to price discriminate.

For future research, I plan to extend my research on the interaction between competition and pricing policies in several respects. First, I plan to improve the quality of my data by collecting more data in locations where markets can be more clearly delineated. This extension would allow me to get better estimates of the price-cost margins for large dealerships like Carmax. Second, I plan to collect cross-sectional data of more dealerships to help further investigate the factors affecting dealerships' choices of pricing policies. Lastly, I plan to explore theoretically the conditions that lead to the asymmetric choices of pricing policies in equilibrium. As my empirical analysis has shown that dealers choosing different pricing policies seem consistent with an asymmetric equilibrium in pricing strategies, a theoretical investigation would shed more lights on the general causes of the observed heterogenous choices of pricing policies.

References

- [1] Armstrong, Mark; Vickers, John. “Competitive Price Discrimination.” *RAND Journal of Economics*, Winter 2001, 32 (4), 579-605.
- [2] Ayres, Ian and Siegelman, Peter. “Race and Gender Discrimination in Bargaining for a New Car.” *American Economic Review*, 1995, 85 (3), 304-321.
- [3] Berry, Steve and Pakes, Ariel. “The Pure Characteristics Demand Model.” *International Economic Review*, 2007, 48 (4),
- [4] Berry, Steven; Levinsohn, James and Pakes, Ariel. “Automobile Prices in Market Equilibrium.” *Econometrica*, July 1995, 63 (4), 841-890.
- [5] Berry, Steven; Levinsohn, James and Pakes, Ariel. “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market.” *The Journal of Political Economy*, February 2004, 112 (1), 68-105.
- [6] Berry, Steven T. “Estimating Discrete-Choice Models of Product Differentiation.” *The RAND Journal of Economics*, Summer 1994, 25 (2), 242-262.
- [7] Bester, Helmut. “Bargaining versus Price Competition in Markets with Quality Uncertainty.” *American Economic Review*, 1993, 83 (1), 278-288.
- [8] Borenstein, Severin. “Price Discrimination in Free-Entry Markets.” *The RAND Journal of Economics*, Autumn 1985, 16 (3), 380-397.
- [9] Goldberg, Pinelopi Koujianou. “Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey.” *Journal of Political Economy*, June 1996, 104 (3), 622-654.
- [10] Goldberg, Pinelopi Koujianou. “Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry.” *Econometrica*, July 1995, 63 (4), 891-951.
- [11] Gul, Faruk and Sonnenschein, Hugo. “On Delay in Bargaining with One-Sided Uncertainty.” *Econometrica*, 1988, 56 (3), 601-611.
- [12] Gul, Faruk and Sonnenschein, Hugo, Wilson, Robert. “Foundations of Dynamic Monopoly and the Coase Conjecture.” *Journal of Economic Theory*, 1986, 39 (1),
- [13] Harless, David W. and Hoffer, George E. “Do Women Pay More for New Vehicles? Evidence from Transaction Price Data.” *American Economic Review*, March 2002, 92 (1), 270-279.

- [14] Holmes, Thomas J. "The Effects of Third-Degree Price Discrimination in Oligopoly." *The American Economic Review*, March 1989, 79 (1), 244-250.
- [15] Kennan, John and Wilson, Robert. "Bargaining with Private Information." *Journal of Economic Literature*, March 1993, 31 (1), 45-104.
- [16] Leslie, Phillip. "Price Discrimination in Broadway Theater." *Rand Journal of Economics*, 2004, 35 (3),
- [17] List, John A. "The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field." *Quarterly Journal of Economics*, 2004, 119 (1),
- [18] List, John A. and Price, Michael K. "Conspiracies and Secret Price Discounts in the Marketplace: Evidence from Field Experiments." *Rand Journal of Economics*, 2005, 36 (3),
- [19] McFadden, D. "Modelling the Choice of Residential Location." *Cowles Foundation Discussion Paper*, 1977,
- [20] McFadden, Daniel. "Econometric Models of Probabilistic Choice," Manski, Charles F., McFadden, Daniel, *Structural Analysis of Discrete Data with Econometric Applications*. the MIT Press, 1981.
- [21] McManus, Brian. "Nonlinear Pricing in an Oligopoly Market: The Case of Specialty Coffee." *Rand Journal of Economics*, 2007, 38 (2),
- [22] Petrin, Amil. "Quantifying the Benefits of New Products: The Case of the Minivan." NBER *Working Paper*, 2001,
- [23] Petrin, Amil. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy*, August 2002, 110 (4), 705-729.
- [24] Riley, John and Zeckhauser, Richard. "Optimal Selling Strategies: When to Haggle, When to Hold Firm." *Quarterly Journal of Economics*, 1983, 98 (2),
- [25] Rochet, Jean-Charles and Stole, Lars A. "Nonlinear Pricing with Random Participation." *Review of Economic Studies*, January 2002, 69 (1), 277-311.
- [26] Salop, Steven and Stiglitz, Joseph. "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion." *Review of Economic Studies*, October 1977, 44 (3), 493-510.
- [27] Salop, Steven C. "Monopolistic Competition with Outside Goods." *Bell Journal of Economics*, 1979, 10 (1), 141-156.
- [28] Scott-Morton, Finon and Zettelmeyer, Florian. "Consumer Information and Discrimination: Does Internet Affect the Pricing of New Cars to Women and Minorities." *Quantity Marketing Economics*, 2003, 1 (1),

- [29] Scott-Morton, Fiona; Zettelmeyer, Florian and Silva-Risso, Jorge. “A Test of Bargaining Theory in the Auto Retailing Industry.” *Yale University Working Paper*, 2004,
- [30] Stole, LA. “Nonlinear Pricing and Oligopoly.” *Journal of Economics & Management Strategy*, 1995, 4 (4),
- [31] Stole, LA. “Price Discrimination and Competition,” *Handbook of Industrial Organization*. 2007, 2221-2299.
- [32] Wang, Ruqu. “Auctions versus Posted-Price Selling.” *American Economic Review*, 1993, 83 (4),
- [33] Wang, Ruqu. “Bargaining versus Posted Price Selling.” *European Economic Review*, 1995, 39 (9),
- [34] Wernerfelt, Birger. “Selling Formats for Search Goods.” *Marketing science*, 1994, 13 (3),
- [35] Wilson, Robert B. *Nonlinear Pricing*. Oxford University Press, 1993.
- [36] Yang, Huanxin and Lixin, Ye. “Nonlinear Pricing, Market Coverage, and Competition.” *Theoretical Economics*, 2008, 3 (1),

Appendix

In this appendix I show the solution to the optimal monopoly pricing problem described in the main text. Let $U(v) = \max_{v'} (v - p(v')) \phi(v')$ be the value function of a buyer with reservation value v . As the buyer's utility function in the problem satisfies the single-crossing property in (v, ϕ) , the incentive compatibility constraint in the optimal pricing problem can be shown to be equivalent to the following two conditions:

$$(i) U(v) = \int_{\underline{v}}^v \phi(t) dt + U(\underline{v})$$

(ii) $\phi(v)$ is increasing in v .

Given the *IR* condition, and $U(v)$ being increasing in v , we must have $U(\underline{v}) = 0$ for the optimal solution. The standard approach to further solve the problem is to first solve the problem ignoring the monotone condition, and then check if the solution found satisfy the monotone condition. Given $U(v) = (v - p(v)) \phi(v)$, the condition (i) is equivalent to

$$p(v) = v - \int_{\underline{v}}^v \phi(t) / \phi(v) dt$$

By substituting the price function $p(v)$ into the seller's payoff function, we have the simplified sellers problem as

$$\max_{\phi(\cdot)} \int_{\underline{v}}^{\bar{v}} \left[(v - c) \phi(v) - \int_{\underline{v}}^v \phi(t) dt \right] dF(v)$$

where $(v - c) \phi(v)$ is the total surplus of selling to a buyer with reservation value v , and $\int_{\underline{v}}^v \phi(t) dt$ is the buyer's surplus from the trade. Through integration by parts, we have $\int_{\underline{v}}^{\bar{v}} \int_{\underline{v}}^v \phi(t) dt dF(v) = \int_{\underline{v}}^{\bar{v}} (1 - F(v)) \phi(v) dv$. Thus the seller's problem further simplifies to be

$$\max_{\phi(\cdot)} \int_{\underline{v}}^{\bar{v}} \left[\left(v - \frac{1 - F(v)}{f(v)} - c \right) \phi(v) \right] dv$$

Suppose that $\frac{d}{dv} \left(\frac{f(v)}{1 - F(v)} \right) \geq 0$,⁵⁹ then $\left(v - \frac{1 - F(v)}{f(v)} \right)$, the "virtual surplus" obtained by the seller conditional the sales probability, is increasing. As $\phi(v) \in [0, 1]$, it is then obvious that the solution to the above problem is

$$\phi^*(v) = \begin{cases} 1, & \text{if } v \geq v^* \\ 0, & \text{if } v < v^* \end{cases}$$

where v^* is the solution to the equation: $v - \frac{1 - F(v)}{f(v)} - c = 0$. And accordingly,

$$p^*(v) = \begin{cases} v^*, & \text{if } v \geq v^* \\ 0, & \text{if } v < v^* \end{cases}$$

As ϕ^* is increasing in v , the monotone condition (ii) is satisfied. The problem can also be solved

⁵⁹This condition is often called the monotone hazard rate condition. This condition is satisfied if F is, for example, uniform, normal, logistic, chi-squared, exponential, or Laplace.

as an optimal control problem by defining $\int_{\underline{v}}^v \phi(t) dt$ as a state variable.

Tables and Figures

Table 1. Data Summary Statistics

| Variables | N | Mean | St. Dev | Max | Min |
|----------------------|------|--------|---------|--------|--------|
| Market Share | 1060 | 0.0039 | 0.0059 | 0.0007 | 0.0543 |
| Price | 1060 | 14.680 | 5.303 | 7.332 | 34.998 |
| Fixed Pricing | 1060 | 0.690 | 0.463 | 0 | 1 |
| Engine Volume | 1060 | 2.849 | 0.951 | 1.3 | 6 |
| Age | 1060 | 2.910 | 1.778 | 0 | 8 |
| All Wheel Drive | 1056 | 0.109 | 0.290 | 0 | 1 |
| Quality Inspection | 1056 | 0.505 | 0.500 | 0 | 1 |
| Manual Transmission | 1056 | 0.151 | 0.305 | 0 | 1 |
| American Car | 1056 | 0.388 | 0.488 | 0 | 1 |
| European Car | 1056 | 0.193 | 0.395 | 0 | 1 |
| Mileage (1000 miles) | 1056 | 41.818 | 19.319 | 0.010 | 98.935 |
| Inventory size* | 1056 | 1.652 | 1.314 | 1 | 14 |

Note: the inventory sizes are measured at the model level.

Table 2. Market Structure Exemplified by Statistics from the Last Quarter

| Statistics | Al Packer | Carmax | Castle | National | Schaefer |
|--|-----------|--------|--------|----------|----------|
| Officially No-Haggling | No | Yes | Yes | No | No |
| # of Models Carried | 18 | 133 | 49 | 42 | 23 |
| Market Shares by Sales Quantity (%) | 2.9 | 77.6 | 8.7 | 6.8 | 3.7 |
| Share of American Cars (%) | 57.6 | 35.3 | 46.2 | 18.1 | 67.8 |
| Share of European Cars (%) | 1.4 | 20.6 | 14.9 | 35.5 | 5.6 |
| Share of other (mostly Asian) Cars (%) | 41.0 | 44.1 | 38.9 | 46.4 | 26.6 |
| Average # of Direct Competitors | 3.19 | 1.87 | 2.74 | 2.70 | 2.94 |

Notes: 1. Carmax Al Packer, Castle, and Schaefer also have new car business.

2. By “direct competitors” I mean other dealerships that also sell the same model; and the average is taken over all models sold by each dealership respectively.

Table 3. Hedonic Regression Results

| Dependent Variable: Price | | |
|------------------------------------|---------------------|----------------|
| Independent Variables | Coefficients | Standard Error |
| Engine Capacity | 2.883** | 0.720 |
| Age | -0.560** | 0.231 |
| Mileage | -0.133** | 0.028 |
| (Asian Car)XMileage | 0.124** | 0.022 |
| ((Asian Car)XMileage) ² | -0.001** | 0.000 |
| All-Wheel Drive | 11.651** | 2.255 |
| (All-Wheel Drive) ² | -8.629** | 2.336 |
| Manual Transmission | 1.871** | 0.365 |
| American Car (dummy) | -0.012 | 0.423 |
| European Car (dummy) | 6.276** | 0.485 |
| Al Packer (dummy) | 0.588 | 0.360 |
| Carmax (dummy) | 2.230** | 0.280 |
| National (dummy) | -0.578 ⁺ | 0.312 |
| Schaefer (dummy) | 1.379** | 0.383 |
| Constant | 8.465** | 1.150 |
| Obs | | 1056 |
| R-squared | | 0.65 |

⁺ significant at the 10% confidence level; * significant at the 5% confidence level; ** significant at the 1% confidence level.

Table 4. The Estimated Demand System Coefficients

| Variables | logit IV | Baseline Model | Augmented Model $\beta = 0.98$ |
|-------------------------------|---------------------|---------------------|-----------------------------------|
| Engine Volume | 0.523** (0.144) | 0.606** (0.883) | 0.834** (0.233) |
| Age | -0.600** (0.214) | -0.604 (0.515) | -1.036** (0.434) |
| (Asian Cars)XMileage | 0.259** (0.052) | 0.278** (0.058) | 0.336** (0.078) |
| All Wheel Drive | 3.159** (0.518) | 2.809** (0.781) | 3.854** (0.688) |
| Quality Inspection | 1.250** (0.090) | 1.226** (0.080) | 1.649** (0.153) |
| Manual Transmission | 0.208** (0.098) | 0.215** (0.105) | 0.477** (0.159) |
| American Car | 0.175* (0.089) | 0.176 (0.114) | 0.096 (0.117) |
| European Car | 0.921** (0.180) | 0.891** (0.171) | 1.387** (0.286) |
| Price | -0.159** (0.029) | | |
| $\log(y_i - p_j)$ | | 7.639** (1.563) | 12.512** (2.788) |
| Discount at Al Packer | | | 1.969** (0.574) |
| Discount at National | | | 0.733 ⁺ (0.444) |
| Discount at Schaefer | | | 2.080** (0.440) |
| σ_1 | | 0.049 (5.048) | 0.001 (5.490) |
| σ_2 | | 1.192 (1.123) | 2.279** (0.769) |
| σ_3 | | 0.047 (1.519) | 0.008 (3.486) |
| $\log(\text{Inventory size})$ | 0.873** (0.048) | 0.918** (0.056) | 0.904** (0.060) |
| $Var(\text{price})$ | 0.004** (0.002) | 0.003 (0.003) | 0.005** (0.002) |
| $Cov(\text{age,price})$ | -0.199** (0.053) | -0.184** (0.086) | -0.228** (0.059) |
| Constant | -8.244** (0.349) | -5.530** (1.056) | -4.616** (0.691) |
| Quarter Dummies | Yes | Yes | Yes |
| Observations | 1060 | 1060 | 1060 |
| R-squared | 0.53 | | |

Note: Standard errors in parentheses; ⁺significant at the 10% confidence level; *significant at the 5% confidence level; ** significant at the 1% confidence level. The prices and discounts are measured in \$1000.

Table 5. Alternative Specification, and the Sensitivity of Results to the Value of β

| Variables | The Step-Two Results of the Two-Step GMM | | | |
|--------------------------------|--|--------------------------------------|--------------------------------------|-------------------------------------|
| | Baseline Model with dealer dummies | Augmented Model $\beta = 0.98$ | Augmented Model $\beta = 0.95$ | Augmented Model $\beta = 0.9$ |
| Engine Volume | 0.845** (0.327) | 0.834** (0.233) | 0.832** (0.335) | 0.831** (0.336) |
| Age | -1.012** (0.446) | -1.036** (0.434) | -1.079** (0.376) | -1.050** (0.467) |
| Asian Cars*Mileage | 0.339** (0.077) | 0.336** (0.078) | 0.335** (0.077) | 0.333** (0.074) |
| All Wheel Drive | 3.927** (0.666) | 3.854** (0.688) | 3.825** (0.731) | 3.735** (0.685) |
| Quality Inspection | 1.669** (0.158) | 1.649** (0.153) | 1.640** (0.218) | 1.614** (0.150) |
| Manual Transmission | 0.480** (0.162) | 0.477** (0.159) | 0.475** (0.157) | 0.467** (0.154) |
| American Car | 0.092 (0.121) | 0.096 (0.117) | 0.098 (0.129) | 0.103 (0.115) |
| European Car | 1.408** (0.288) | 1.387** (0.286) | 1.379** (0.269) | 1.351** (0.268) |
| $\log(y_i - p_j)$ | 12.728** (2.936) | 12.512** (2.788) | 12.425** (2.617) | 12.148** (2.477) |
| Dummy/Discount of Al Packer | 0.492** (0.120) | 1.969** (0.574) | 2.498** (0.661) | 3.344** (0.889) |
| Dummy/Discount of National | 0.106 (0.098) | 0.733 (0.444) | 1.143** (0.595) | 1.721** (0.852) |
| Dummy/Discount of Schaefer | 0.483** (0.140) | 2.080** (0.440) | 2.603** (0.517) | 3.410** (0.733) |
| σ_1 (engine capacity) | -0.017 (5.100) | 0.001 (5.490) | 0.006 (5.444) | 0.016 (5.391) |
| σ_2 (constant) | 2.401** (0.910) | 2.279** (0.769) | 2.230** (0.815) | 2.137** (0.948) |
| σ_3 (age) | 0.009 (3.645) | 0.008 (3.486) | 0.008 (3.446) | 0.010 (3.482) |
| $\log(\text{Inventory size})$ | 0.899** (0.059) | 0.904** (0.060) | 0.906** (0.062) | 0.916** (0.060) |
| $Var(\text{price})$ | 0.005** (0.002) | 0.005** (0.002) | 0.005** (0.002) | 0.005** (0.002) |
| $Cov(\text{age,price})$ | -0.232** (0.059) | -0.228** (0.059) | -0.221** (0.067) | -0.217** (0.059) |
| Constant | -4.559** (0.841) | -4.616** (0.691) | -4.642** (1.503) | -4.713** (0.804) |
| Quarter Dummies | Yes | Yes | Yes | Yes |
| Observations | 1060 | 1060 | 1060 | 1060 |

Notes: Standard errors in parentheses. + significant at the 10% confidence level;

* significant at the 5% confidence level;** significant at the 1% confidence level.

The prices and discounts are measured in \$1000.

Table 6. The Discounts Obtained in the Experiment in Ayres and Siegelman (1995)

| Consumer subgroups | Number of Observations | Initial Discount | Final Discount |
|--------------------|------------------------|------------------|----------------|
| White Male | 153 | 1068 | 1476 |
| White Female | 53 | 876 | 1301 |
| Black Female | 60 | 664 | 971 |
| Black Male | 40 | 0 | 233 |

The weighted average final discount: \$1184

Source of Information: Table 2 in Ayres and Siegelman (1995).

Note: In computing the above numbers, I used the profit at the initial offer for black male testers as the profit at the list price. The initial discounts and final discounts are respectively the discounts at the initial offer prices and the final offer prices. The testers waited for about 5 minutes to get the initial offer prices.

Table 7. The Comparison of the Model-Predicted and KBB Margins

| | Al Pakcer | Carmax | Castle | National | Schaefer |
|-----------------|-----------|--------|--------|----------|----------|
| Model-Predicted | 3.34 | 8.59 | 3.18 | 3.74 | 3.24 |
| KBB Compiled | 3.40 | 3.74 | 3.60 | 3.80 | 3.33 |

Note: The margins are the market-share-weighted average margins for the dealerships. The KBB price-cost margins are computed by me based on information provided by KBB. Both margins are net of discounts and measured in \$1000. For the KBB margins I assumed that the market average discount is \$500.

Table 8.1. Al Packer Comparison Results

| | Discounts | Total Profits | Sales Quantity | Average Margin |
|--------------------|-----------|---------------|----------------|----------------|
| Posted Price | 0 | 17.501 | 100% | 4.29 |
| Optimal Screening | 1.151 | 18.325 | 104.7% | 3.54 |
| Actual (estimated) | 1.969 | 15.15 | 86.6% | 3.34 |

Note: Profits and average margins are measured in \$1000. Al Paker haggles.

Table 8.2. Carmax Comparison Results

| | Discounts | Total Profits | Sales Quantity | Average Margin |
|-------------------|-----------|---------------|----------------|----------------|
| Posted Price | 0 | 2203.3 | 100% | 8.59 |
| Optimal Screening | 0.602 | 2227.10 | 101.1% | 7.85 |

Note: Profits and average margins are measured in \$1000. Carmax does not haggle.

Table 8.3. Castle Comparison Results

| | Discounts | Total Profits | Sales Quantity | Average Margin |
|-------------------|-----------|---------------|----------------|----------------|
| Posted Price | 0 | 105.79 | 100% | 3.18 |
| Optimal Screening | 0.761 | 108.10 | 102.2% | 2.75 |

Note: Profits and average margins are measured in \$1000. Castle does not haggle.

Table 8.4. National Comparison Results

| | Discounts | Total Profits | Sales Quantity | | Average Margin | | |
|--------------------|-----------|---------------|----------------|-------|----------------|------|-------|
| Posted Price | 0 | 44.20 | 100% | 11.65 | 100% | 3.81 | 100% |
| Optimal Screening | .625 | 46.37 | 104.9% | 12.06 | 103.5% | 3.80 | 99.7% |
| Actual (estimated) | 0.733 | 42.00 | 95.0% | 11.22 | 96.3% | 3.74 | 98.1% |

Note: Profits and average margins are measured in \$1000. National haggles.

Table 8.5. Schaefer Comparison Results

| | Discounts | Total Profits | Sales Quantity | | Average Margin | | |
|--------------------|-----------|---------------|----------------|------|----------------|------|-------|
| Posted Price | 0 | 24.96 | 100% | 6.33 | 100% | 4.29 | 100% |
| Optimal Screening | 1.14 | 26.73 | 107.1% | 7.07 | 111.9% | 3.72 | 86.7% |
| Actual (estimated) | 2.08 | 21.39 | 80.0% | 6.77 | 107.0% | 3.24 | 75.5% |

Note: Profits and average margins are measured in \$1000. Schaefer haggles.

Table 9 Market Structure and Returns to Switching from Posted Price to Haggling

| Statistics | Al Packer | Carmax | Castle | National | Schaefer |
|--------------------------------------|-----------|--------|--------|----------|----------|
| Officially No-Haggling | No | Yes | Yes | No | No |
| # of Models Carried | 18 | 133 | 49 | 42 | 23 |
| Market Shares by Sales Quantity (%) | 2.9 | 77.6 | 8.7 | 6.8 | 3.7 |
| Average # of Direct Competitors | 3.19 | 1.87 | 2.74 | 2.70 | 2.94 |
| Returns to Switching to Haggling (%) | 4.7 | 1.08 | 2.2 | 4.9 | 7.1 |

Notes: 1. Carmax Al Packer, Castle, and Schaefer also have new car business.

2. By “direct competitors” I mean other dealerships that also sell the same model; and the average is taken over all models sold by each dealership respectively.

Table 10.1 The Increases in Profit When Baby Carmaxes Supplying Cars with Increasingly Higher Quality Switch from Posted Price to Hagging

| Baby Carmaxes | Profit Increase | Optimal Discount |
|---------------|-----------------|------------------|
| 1 | 2.12 | 0.7378 |
| 2 | 2.76 | 0.7361 |
| 3 | 4.05 | 0.6759 |
| 4 | 3.77 | 0.6265 |
| 5 | 4.77 | 0.6250 |
| 6 | 4.54 | 0.6251 |
| 7 | 3.81 | 0.6254 |
| 8 | 3.38 | 0.6245 |
| 9 | 2.96 | 0.6250 |
| 10 | 3.25 | 0.6248 |
| 11 | 2.53 | 0.6249 |
| 12 | 2.02 | 0.6247 |

Note: 1. The 12 “Baby Carmaxes” are counterfactual dealerships I created using the pool of models supplied by Carmax. The Baby Carmaxes in this set of experiments are formed based on the mean qualities of the models: the n th Baby Carmax carries models with mean qualities in the range from the $((n - 1) \cdot 100/12)$ th percentile to the $(n \cdot 100/12)$ th percentile of the empirical distribution of the mean qualities of the models supplied by Carmax.
 2. The profit increases are measured in percentage.
 3. The optimal discounts are presented in \$1000.

Table 10.2 The Increases in Profit When Baby Carmaxes Supplying Different Number of Models Switch from Posted Price to Hagging

| | Baby Carmaxes Defined by the Number of Models Carried by Them | | | |
|--------------------------|---|--------------------|--------------------|--------------------|
| | 43 Models | 21 Models | 15 Models | 11 Models |
| Average Profit Increase | 1.98 (0.94) | 2.70 (1.76) | 2.89 (1.61) | 3.62 (1.86) |
| Average Optimal Discount | 0.6199 (0.0496) | 0.5972 (0.0946) | 0.6111 (0.0761) | 0.5953 (0.1013) |
| Number of Experiments | 50 | 50 | 50 | 50 |

Note: 1. The “Baby Carmaxes” are counterfactual dealerships I created based on the pool of models carried by Carmax. The Baby Carmaxes in this set of experiments are created by randomly sampling without replacement the respectively number of models from the pool of models supplied by Carmax.
 2. The proportional profit increases are measured in percentage; and the optimal discounts are measured in \$1000.
 3. The standard deviations of the proportional profit increases and the optimal discounts are in the parenthesis.

Figure 1.1 Snapshot of the Main Page Carmax's Website

CARmax

Find cars by... **SMART**

Make
Vehicle Type
Model & Price
Features
MPG

Affordable cars with low, no-haggle prices

Thinking about selling your car? Come to CarMax.

You'll get a **free appraisal** while you wait
Receive a written offer good for 7 days
We'll buy your car even if you don't buy ours!®
[Contact us about appraisals.](#)

100 locations across the country
There's a CarMax near you.

Need financing? We can help.
With approved credit, CarMax works with multiple financial institutions to offer competitive financing for almost every budget.
[Learn more about financing at CarMax.](#)

Enjoy peace of mind
5-Day Money-Back Guarantee*
Free Vehicle History Report
Limited 30-Day Warranty (60-Day in CT)*
Reliability ratings from J.D. Power
NHTSA safety ratings
*See store for written details

Now hiring - Sales Consultants
Looking for a rewarding and challenging career? Now is a great time to work at CarMax! We're hiring at stores across the country.
[Apply now](#)

Note: the “no-haggle” pricing policy is explicitly stated in the upper left of the webpage.

Figure 1.2 Snapshot of the Main Page the Caslte Auto Group's Website

CASTLE AUTO GROUP
"Where the customer is King" "No hassle at Castle"

888-503-4411
CASTLE FORD
3111 Automobile Blvd. Silver Spring, MD

888-503-4411
CASTLE MITSUBISHI
3111 Automobile Blvd. Silver Spring, MD

888-503-4411
CASTLE MAZDA
3111 Automobile Blvd. Silver Spring, MD


800-825-7941
CASTLE AUTO OUTLET
921 Pulaski Highway Joppa, MD 21085

410-803-0800
CASTLE MOTORSPORTS
11717 Belair Road Kingsville, MD 21087

410-803-0800
CASTLE MARINE
11717 Belair Road Kingsville, MD 21087

Note: the “no-hassle” pricing policy is officially stated as well.

Figure 2.1 the Identification of d_j Given that β is Known



Note: $v(d_j | \beta, \alpha, y_i, p_j)$ is the net value of the discount d_j to a consumer with

$\delta_j + \eta_{ij} + \alpha \log\left(\frac{y_i - p_j}{y_i}\right) = 0$. Product k is consumer i 's most preferred choice among all products excluding j . The model predicted probability of consumer i purchasing product j is

$$\int \int \int_{\eta_{ij} > \eta_{ij}^*} S_{ij}(\eta_{ij} | y_i, \epsilon_{ij}, \epsilon_{ik}, \eta_{ik}) dP(\eta_{ij} | \eta_{ik}) dP(\eta_{ik}) dP((y_i, \epsilon_{ij}, \epsilon_{ik}))$$

Figure 2.2 the Lack of Joint Identification of (β, d_j)



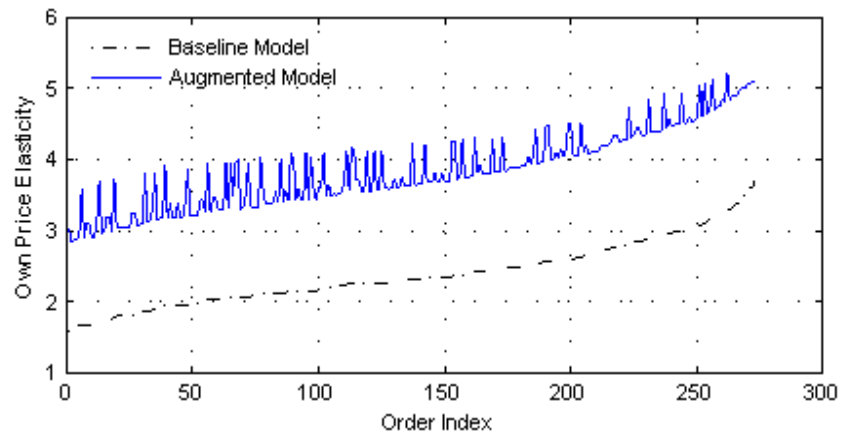


Figure 2.3 the Underestimate of α by the Baseline Model



Note: in the figure, $\alpha' < \alpha$ and $p_j > p_k$.

Figure 3. Own Price Elasticities Estimated in the Baseline and Augmented Model



Note: the two sets of own price elasticities are sorted according to the one estimated under the baseline model; the elasticities presented in the figure are the ones computed for the latest quarter covered by my data.

Figure 4.1 The Information about the Chrysler Pt Cruiser on the Dealership's Website



Al Packer's Whitmarsh Ford
 9801 Pulaski Highway
 Baltimore Maryland 21220
www.alpackerford.com

PREOWNED INVENTORY

| Filter | Make | Model | Trim | Year | Body Type |
|---------|----------|------------|------|------|-----------|
| Search: | Chrysler | Pt Cruiser | All | All | All |

Search by Stock/VIN Number

[SEARCH AGAIN] Results 1 to 1 of 1

Sort By: Year ▲/▼ | Make ▲/▼ | Mileage ▲/▼ | Price ▲/▼

Page 1 of 1

Select Page: **1**
[Show All Vehicles](#)

2005 Chrysler PT Cruiser

Market Value: \$12,271.00
Packer Price: \$10,995.00



VIN Number: 3C4FY48B95T535497
 Stock Number: WF90417A
 Exterior Color: Midnight Blue Pearlcoat
 Mileage: 63829
 Transmission: Automatic
 Engine: 4 Cyl. 2.4
 Body Type: 4 Dr Wagon

Notes:
Stop by or call and talk with one of our trained Used Car specialists today. Dont see what you want... we can get it for you TODAY! Check out our GREAT SPECIALS at www.alpacker.com Prices dont include taxes, registration or \$100 Dealer processing fee



Details and Media
 View Window Sticker
 Contact Us
 Make Offer

[SEARCH AGAIN] Results 1 to 1 of 1

Page 1 of 1

Select Page: **1**


Our Dealership has posted this vehicle listing solely for the convenience of their customers. Dealer Specialties and Al Packer's Whitmarsh Ford make no representations, expressed or implied, to any actual or prospective purchaser or owner of this vehicle as to existence, ownership, accuracy, description, or condition of this vehicle, listed equipment, accessories, price, or w warranties. Any and all differences must be addressed PRIOR to the sale of the vehicle, Dealer Processing Fee (Not Req'd by Law) \$100.00.

Note: The figure is a snapshot of the Pt Cruiser's listing page we took in the morning of Oct 10, 2009, the day when we went to the dealership to negotiate a price for the car.

Figure 4.2 The 2005 Chrysler Pt Cruiser on the Dealership's Parking Lot




Figure 4.3 The Price Offer We Obtained from Al Packer for the Pt Cruiser



Note: the price offer written on the back of the salesperson's business card: "\$9995+tax/tags freight, Pt Cruiser". The sticker (list) price was \$10995. The discount we got was \$1000.

Figure 4.4 The Price Offer We Obtained from Al Packer for the Pt Cruiser



Note: the above scanned images are the two sides of the business card of the salesperson we negotiated with. The offer of \$9995 was written by Mr Shea, the salesperson, on the flip side of the card.

Figure 5.1 The Empirical CDF of the Mean Qualities of Cars Sold by the Dealers

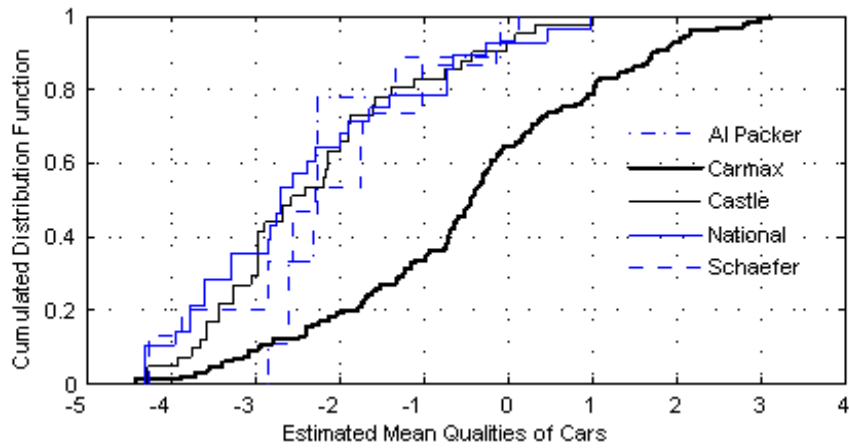
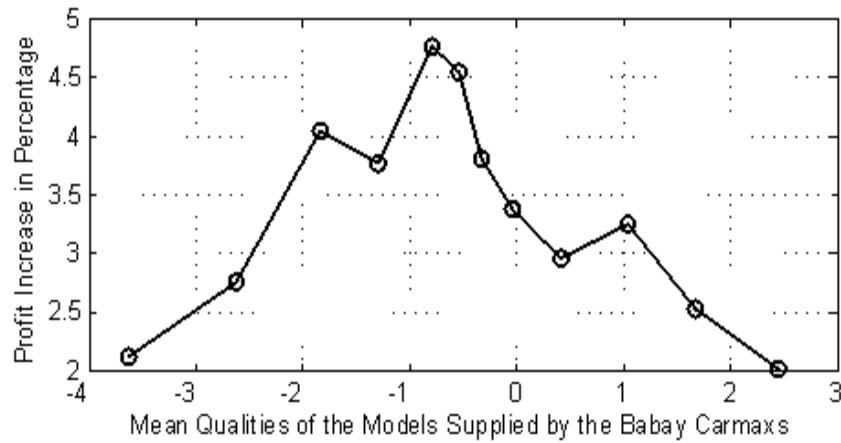


Figure 5.2 The Profit Increase if Baby Carmaxes Switched to Hagglng



Note: Each point represents a Baby Carmax. The Baby Carmaxes in this graph are formed based on the mean qualities of the models supplied by Carmax.