

Invest in Information or Wing It?

A Model of Dynamic Pricing with Seller Learning*

Guofang Huang
Yale School of Management
Hong Luo
Harvard Business School
Jing Xia
Harvard University

June, 2012

Abstract

This paper studies the managerial problem of dynamic pricing in the secondary durable-goods market, where sellers typically have limited information about item-specific heterogeneity. It develops a structural model of dynamic pricing that features the seller learning about item-specific demand through initial assessment and active learning in the sale process. The model is estimated using novel panel data of a leading used-car dealership. Policy experiments are conducted to quantify the value of the dealer's initial information about item-specific demand and of lowering the price-adjustment cost. With the dealer's average net profit per car in the estimation sample being around \$740, the initial information about item-specific demand worth roughly \$243, and cutting the dealer's price-adjustment cost by half would increase its profit by about \$103.

Keywords: Dynamic Pricing, Item-specific Demand, Active Learning, the Value of Information.

*Guofang Huang: guofang.huang@yale.edu; Hong Luo: hluo@hbs.edu; Jing Xia: jingxia@fas.harvard.edu. We thank Cars.com for providing the data used in this paper, and especially thank Neil Feuling and Anoop Tiwari from Cars.com for their interest in the project and assistance with data construction. We also thank Dirk Bergemann, Guido Imbens, Przemek Jezioriski, Zhenyu Lai, Ariel Pakes, Greg Lewis, Jiwoong Shin, K. Sudhir, Juuso Valimaki and Zhixiang Zhang for helpful discussions. Any remaining errors are our own.

1 Introduction

It is a challenging task for managers to set prices for products in secondary markets, such as used cars, houses, art, etc.. These products show significant item-specific heterogeneity even after accounting for all their standard observable attributes. Take the example of used cars. Identical new cars can end up as used cars with the same mileage but in quite different conditions, depending on how their prior owners have driven and maintained them.¹

As far as the heterogeneity matters to consumers, sellers could incorporate such information into prices to achieve higher profits. One can potentially gain information on item-specific heterogeneity through a few channels. Used-car dealerships, for example, may carefully inspect the cars they acquired before selling them. Furthermore, they can also learn about car-specific heterogeneity in the sale process, for example, through observing the events of no sale or communicating with buyers after they have inspected and test-driven the cars.

Two types of price-setting practices are common in the market when sellers face significant item-specific heterogeneity. One is to set prices by adding a markup, based simply on experience, to the acquisition cost. The other is to assess item-specific heterogeneity carefully and set prices contingent on the assessment results. CarMax, the largest used-car dealership in the U.S., is a nice example of the second type of practice. It puts every used car it acquires through a thorough inspection process before putting them up for sale. The inspection provides the company with information on the conditions of individual cars beyond what can be estimated by the cars' model year and mileage. Furthermore, the dealer's proprietary information management and pricing system enables it to set and adjust prices based on the information it acquires and learns in every stage of the sale process.

Though potentially beneficial, the inspection and assessments require investments in equipment, information-management system and nontrivial variable costs. Should firms go through all the trouble to carefully examine each individual item and set prices accordingly? More specifically, what is the value of the information sellers can obtain in the initial assessment process? To answer these questions, we develop a structural model of dynamic pricing in the presence of seller learning, and estimate it using novel panel sales data of CarMax. We use the estimated structural model to quantify the value of the information generated by the initial assessment and of lowering the price-adjustment cost.

Our theoretical model of dynamic pricing is cast as a stochastic optimal-stopping-time problem. More specifically, we describe item-specific heterogeneity by a scalar

¹As an example of the economic value of the car-specific conditions, based on the information from kbb.com in May 2012, for a 2007 Honda Accord LX sedan with 68,500 miles, the Kelley Blue Book "private party" price is \$12,550 for "excellent condition" and \$11,700 for "good condition" (see Figure 1).

random variable, ξ . Before selling an item, the seller receives a signal—which quantifies the result of the seller’s initial assessment—about ξ from a distribution centered around ξ . We adopt the Bayesian Normal learning framework to capture the seller’s learning process, where we assume that every buyer reveals to the seller his idiosyncratic value for the item—also signals about ξ —if he decides not to buy the item. The seller incurs a fixed cost whenever she changes prices. The inventory evolves as an exogenous first-order Markov process. The seller’s objective is to choose prices over time to maximize the present value of her expected profit from each given item. The problem shares features of the optimal-stopping-time problem in that the prices set by the seller control the probability of sale (stopping) in each stage.

The demand side is captured by a simple static discrete choice model with item-specific unobserved heterogeneity ξ . Given our seller-side model, prices are directly correlated with ξ , which needs be dealt with in estimation.

We derive couple of insights about the optimal pricing strategy from the model. First, the learning dynamics generates a dynamic continuation-value effect on price. Compared to the maximum static profit, the continuation value drops as the value of new information to be learned in the sale process diminishes over time. Therefore, at the beginning, to obtain the significant benefit of learning, the seller would strategically increase the probability of continuing (as opposed to stopping) by setting higher initial prices. In later days, such incentives gradually disappear, and, as a result, prices drop even if the seller’s expectation about item-specific demand does not change. Second, there is an additional controlled-learning effect on price dynamics. The current price set by the seller controls the distribution of the signals she receives conditional on continuing. Lower price implies that the signal distribution conditional on continuing will be more concentrated on lower values, which leads to smaller gains from learning. Therefore, the seller has an incentive to increase prices in order to take full advantage of the opportunity of learning. This second effect is also stronger at the beginning, when learning is more important, than in later days. Together, these two effects imply that the optimal prices tend to drop more, on average, over time than what would have been caused only by selection on ξ and adaptive learning.

We apply our model to a 2011 car-level panel sales data for CarMax. The data include detailed car attributes and all the prices set by the dealer over each car’s entire duration on the market.² Following Petrin and Train (2010), we estimate the demand model by using the control function approach. The method controls for the endogeneity in price in discrete choice models, assuming the existence of an excluded variable in the reduced-form pricing equation and assuming that the shocks affecting prices are normal random variables. We estimate the structural model of dynamic pricing using the Nested Fixed Point algorithm, following Rust (1987). The difficulty in the estimation of our dynamic pricing model lies in the fact that the state variables summarizing the

²The prices at the dealer are nonnegotiable.

seller's belief about ξ are not observable to us and need to be integrated out from the likelihood function. To deal with the difficulty of high dimensional integration over serially correlated random variables, we compute the observable likelihood by simulation, using the method of Sampling and Importance Re-sampling.

Our policy experiments reveal sizable value of the information that CarMax obtains in the initial assessment: The assessment increases the expected profit by around \$243 per car in our sample. The value is significant given the reported net profit of around \$740 per car at CarMax. The estimated menu cost is about \$315, which, in our view, captures mainly the cost of assessing new information and coming up with updated optimal prices. The expected profit would increase by about \$103 if the menu cost were cut by half, and would decrease by about \$50 if the menu cost were doubled.

In the U.S., approximately 37 million used cars were sold in 2011, compared to about 11.6 million new cars sold in the same year. The problem we focus on is important, given the importance of the market and the relatively few relevant empirical studies in the literature. To the best of our knowledge, our paper is the first one to estimate a structural model of dynamic pricing in the presence of seller learning, and to use it to quantify the value of information on item-specific demand. Our modeling framework can potentially be applied to analyzing similar problems in other markets. Our insights on price dynamics in the presence of seller learning is also partly new in the related theory literature.

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3 introduces our data and presents some model-free evidence of demand uncertainty and learning. Section 4 sets up the structural demand model and the seller's dynamic pricing model. Section 5 discusses the estimation method. Section 6 presents the empirical results. Section 7 concludes.

2 Related Literature

This paper is closely related to the economic theory literature on dynamic pricing with demand uncertainty (c.f. Rothschild (1974), Grossman, Kihlstrom, and Mirman (1977), Easley and Keifer (1988), Aghion, Bolton, Harris, and Julien (1991), Mirman, Samuelson, and Urbano (1993), Trefler (1993) and Mason and Välimäki (2011)). Most closely related to our work is Mason and Välimäki (2011).³ In their model, the seller learns the state of item-specific demand from observing the events of no sale, which creates an incentive for the seller to raise the current price to increase the benefit from learning. The seller's uncertainty about demand in their model is motivated by the seller's lack

³ There is a large literature in operations research that studies dynamic pricing with uncertainty in demand. See, for example, Xu and Hopp (2005), Aviv and Pazgal (2005) and Araman and Caldentey (2009). However, these papers focus on sales of standardized products and use quite different approaches than ours.

of information on the buyer’s arrival process, and the state of demand is modeled as binary (either high or low). In contrast, in our model, the seller’s uncertainty about demand derives from her lack of detailed information about item-specific heterogeneity, which is modeled as a continuous variable. Therefore, the model of Mason and Välimäki (2011) model may better capture the problem facing individual sellers, who know the item (e.g., the car or house that they own) that they are selling very well, but are not sure how many buyers might potentially be interested in it. Our model, in contrast, may better capture the problem facing, for example, large used-car dealerships or banks selling a large number of foreclosed houses, where the main issue is the sellers’ lack of information on the conditions of individual cars or houses.

The empirical literature on the seller’s dynamic pricing strategy in the presence of dynamics in demand is small. A well-known example is Nair (2007), who investigates the impact of consumers’ forward-looking behavior on seller’s optimal pricing strategy. In his model, the firm has perfect information about demand and sets prices over time to sell standardized products to a population of consumers with heterogeneous price elasticities. The paper’s simulation results show that consumers’ forward-looking behavior significantly limits the seller’s ability to price discriminate inter-temporally. Our paper is different from that of Nair (2007) both in the substantive focus and in the model. Our substantive focus is on quantifying the value of information about item-specific demand in the context of dynamic pricing. From the modeling perspective, in our model, the seller sells a single item to a sequence of buyers with idiosyncratic preferences for the item. In our model, the dynamics in demand come from seller learning in the sale process, and the dynamics in the optimal pricing strategy in our model also come from quite different sources.

We found few empirical works that analyze the impact of demand uncertainty and the value of information about demand in the context of firms’ pricing or other dynamic decisions. A notable exception is Gardete (2012), who investigates the value of the information about market demand in the DRAM market. The author points out that the value of information about demand in the DRAM market derives from the fact that production and capacity adjustment takes a nontrivial period of time, and demand is volatile even in the short term. In his setting, the value of market information to individual sellers is an empirical question, because the value of shared market information in the oligopoly market is theoretically ambiguous. In comparison, in our model, the value of the information about item-specific demand for the monopolist seller is always positive. The complications in quantifying the value of information in our case arises from the additional dynamics in prices induced by the learning process and the seller’s forward-looking behavior.

Other empirical work on pricing strategy in the auto retail market includes Sudhir (2001) and Huang (2009). These papers focus primarily on the strategic considerations in price setting in oligopoly markets. From the modeling perspective, most of these

papers consider pricing strategies in static models with differentiated products and Bertrand competition.

The general Bayesian learning framework used in our model has been widely adopted in the large literature in marketing and industrial organization that studies consumers' learning behavior and its implications for demand in various markets. Well-known examples from this literature include Erdem and Keane (1996), Akerberg (2003), Crawford and Shum (2005) and Erdem, Keane, and Sun (2008).⁴ Our focus on the supply-side pricing decisions clearly sets our paper apart from these papers. In addition, two features of our model are not considered in these papers. First, in our model, the new information that the seller learns each stage is endogenous to the seller's (pricing) decision in the previous stage. As mentioned above, this creates a controlled-learning effect on price. Second, our model allows for selection on unobserved item-specific heterogeneity by the stopping (sale) event, which is a major force driving the price dynamics in our model.

3 Data and Model-Free Analysis

3.1 Data

The data used in this paper are provided by *Cars.com*, the second-largest automotive classified site in the country. Dealers and individual sellers advertise the prices and other detailed information about their cars on the website. *Cars.com* charges individual sellers about \$30 per listing per month and dealers a smaller amount, depending on their listing volume. Although the website lists both new cars and used cars, we use only data on the used-car segment in our analysis. Our data include the entire population of cars listed by dealers within a 20-mile radius of four zip codes between January 2008 and December 2011.⁵ The data contain detailed information on car characteristics and daily list prices of each car for its entire duration on the market.

Typically, a car is removed from the website once the dealership has sold it. There are, however, other possibilities. For example, a car might be taken to a whole sale auction, which is most likely to happen when the car has stayed on the market for too long. Industry reports suggest that such a possibility is very small. For example, CarMax says, in its 2011 annual report, that "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." Another possibility is that a chain dealer might transfer a car to its store at another location in the region. This also seems to be a small problem. In our analysis, we eliminate cars that were ever listed in more than

⁴Ching, Erdem, and Keane (2011) provide a comprehensive survey of the empirical literature of consumer learning.

⁵The four zip codes are 21162, 22911, 64055, and 66202.

one locations.

As we mentioned in the Introduction, CarMax is a particularly good example of systematically acquiring and utilizing information in its pricing decisions. In the local markets that CarMax has entered, it has always been the dominant player. In addition, CarMax lists its entire inventory on *Cars.com*,⁶ and CarMax’s “no-haggle” pricing policy means that the list prices we observe would also be the actual transaction prices. These features of CarMax make it an ideal seller to focus on in our analysis. Therefore, we will use CarMax’s data in our exercise to measure the value of information and learning.

In the rest of this section, we present some preliminary findings about dealers’ pricing behavior using the 2011 data of the top six dealerships, including CarMax, in the area of White Marsh in Baltimore County, Maryland.

3.2 Model-Free Analysis

Evidence of demand uncertainty and seller learning

The strength of our data is that we observe each car’s daily list prices for its entire duration on the market. The patterns in price dynamics that we identified from the data provide some preliminary evidence for the importance of car-specific demand uncertainty and seller learning in this market.

First, a car usually takes quite a few days to sell, which creates the room for the seller to learn new information and to adjust prices. Table 1 shows the summary statistics of the distribution of the time-to-sell (TOS) for the entire sample and for CarMax. Typically, a car stays on the market for 13 days. The distribution of TOS has a relatively long and thick tail on the right. About 25% of all cars take longer than 29 days to sell. Relative to the entire sample, CarMax takes fewer days to sell its cars.

Second, a significant share of cars experience price changes, and the magnitudes of the changes are substantial. Table 2 tabulates the total number of price adjustments during the cars’ entire time on the market. About 36% of all cars have their prices changed at least once, and the maximum number of changes is 11. A smaller share (30%) of cars at CarMax have ever changed price, and the maximum number of changes is also smaller compared to the entire sample.

Given the relatively stable demand for used cars in 2011, and given that most cars are sold within a relatively short time, the price changes here are most likely driven by two main factors. One is inventory shocks, and the other is sellers updating their beliefs about car-specific demand. The evidence we present in the following suggests that the latter might be the main driver of the observed price changes.

Table 3a shows the distribution of one-time price changes—i.e., the change in the

⁶In its 2011 annual report, CarMax says that it “lists every retail used vehicle on both *Autotrader.com* and *Cars.com*.”

current price relative to that of the previous day, conditional on the change being strictly positive. A majority of the one-time price changes are price decreases (89%), though there are also a fair amount of price increases (11%). The magnitudes of one-time price changes are in similar ranges for increases and decreases. The mean one-time price changes are \$924 and \$688, respectively, for price increases and decreases. Table 3b shows the distribution of total price changes—i.e., the change in the cars’ last prices relative to their initial prices, conditional on the total price change being strictly positive. Here, we also see similar patterns in price changes.

The price changes being predominantly decreases suggests that the main driver of price changes is more likely seller learning than inventory shocks. The inventory shock is equally likely to generate price increases and price decreases. But when there is car-specific demand uncertainty and seller learning, the selection on car-specific heterogeneity by time-on-market (TOM) would generate significantly more price decreases than price increases.

Third, there are clear patterns in the price-changing frequency by time, which are consistent with nontrivial menu costs and their interaction with the learning dynamics. We have seen that price adjustments are infrequent but sizable, which suggests that the menu cost could be an important factor affecting dealers’ pricing decisions. Figure 2 plots the share of cars with prices changed relative to their prices on the previous day for each day of TOM. Two patterns stand out in the graph. First, the likelihood of price changes drops quickly over the first few days, and then largely flattens out. Second, the likelihood of price change spikes once before it drops again and flattens out. Both patterns are consistent with the presence of seller learning interacting with the menu cost. In Figure 4, we also see very similar patterns in the graph for the sample of CarMax.

Lastly, Table 4 shows the summary statistics of the daily percentage change of inventories for the entire sample, for CarMax, and for the top six brands of CarMax. We see that for most days, the change in inventories is relatively small, which makes it less likely to be the main reason driving the large changes in price.

Significant variations across dealers

There is significant variation in performance across dealerships. Table 5a presents the summary statistics of the distribution of the TOS by dealer for the top six dealers in our sample. Here, we see large variation across dealers. For example, CarMax, typically takes nine days to sell a car, while Cook and MileOne typically take as long as 32 days. Table 5b summarizes the distribution of total price changes by dealership. The variation across dealerships is, again, quite significant. The mean total price change for CarMax is smaller than that of all the other dealerships. Variation in local demand and inventory composition might partly explain the differences, but we think that the differences in

sellers' information about car-specific demand and in their ability to dynamically adjust prices to incorporate new information arising from the sale process could be the more important factors driving the differences in performance across dealerships.

Table 6 summarizes the same distributions for the top six brands within CarMax. In contrast, we see that the variation is much smaller for both TOS and total price changes. It shows that CarMax's superior performance is consistent across all cars it sells.

Alternative hypothesis

One might argue that if consumers are forward-looking when making used-car purchase decisions, and if they have heterogeneous tastes for price discounts, the seller may also find it optimal to lower price sequentially—that is, to skim the price-inelastic consumers first and then sell to the price-elastic consumers later. However, this does not seem to be a compelling story for the secondary durable-goods market given the patterns of price dynamics found in our data. Furthermore, each used car is a unique product, and normally gets sold within a few days. And, most cars are sold without having their prices ever changed. Therefore, it seems that there is little scope for the forward-looking behavior of consumers in this particular market.

In summary, given the evidences shown above, we find it essential to incorporate item-specific demand uncertainty and seller learning to explain the pricing behavior observed in the data. In the following, we describe a structural model of dynamic pricing with these features. We will apply the model to the used-car sales data described above to analyze the value of information and learning in the problem of price-setting.

4 Model

4.1 Overview

Our model is set up in the context of a used-car retail market. A used-car dealer is setting prices for a car for buyers who sequentially arrive at her store. Each buyer has a demand for, at most, one car. After a buyer makes a one-shot purchase decision, he exits the market. Each buyer tells the seller his value for the car if he decides not to buy the car. The seller inspects the car and evaluates its condition before she starts to sell it. Based on her assessment, the seller sets a price for the car for the first buyer. Until the car is sold, the seller sets a price for it before the arrival of each of the subsequent buyers based on her latest belief about the demand for the car. The price set by the dealer is nonnegotiable.

4.2 Demand

Let the car that we are focusing on be car j . X_{1j} is a vector of observed attributes of car j ; scalar ξ_{1j} summarizes car j 's condition, which is observable to all buyers but not to the econometrician. Suppose that car j 's list price for buyer i is p_{ij} . Let v_{ij} be buyer i 's value for car j . We assume that v_{ij} is determined as follows:

$$v_{ij} = X_{1j}\beta + p_{ij}\alpha + \xi_j + \varepsilon_{ij}$$

where ε_{ij} is buyer i 's idiosyncratic preference shock for car j , and β and α are, respectively, the marginal value for X_{1j} and the price. Furthermore, let buyer i 's value of buying cars other than car j be v_{ia} , and the value of not buying any car be v_{i0} . We specify the values of the two choices as follows:

$$\begin{aligned} v_{ia} &= u_a + \varepsilon_{ia} \\ v_{i0} &= \varepsilon_{i0} \end{aligned}$$

where u_a is the mean value of buying other cars; the mean value of not buying any car is normalized to zero; ε_{ia} and ε_{i0} are buyer i 's idiosyncratic preference shocks for the two choices. We will describe how we approximate the mean value of buying other cars, u_a , in the estimation section.

Define I_{ij} as an indicator function of buyer i choosing to buy car j . Then, we have:

$$I_{ij} = 1 \{v_{ij} > v_{ia} \ \& \ v_{ij} > v_{i0}\}$$

That is, buyer i buys car j if and only if the choice gives him the highest value.

We further assume that the preference shocks $(\varepsilon_{ij}, \varepsilon_{ia}, \varepsilon_{i0})$ are mutually independent, and independent of ξ_j , and that $\varepsilon_{ij}, \varepsilon_{ia}, \varepsilon_{i0}$ are all random variables from the standard normal distribution.

4.3 Dynamic Pricing

For car j , the seller observes X_{1j} , but not ξ_j . Regarding ξ_j , the seller knows only the population distribution of ξ_j before she inspects the car. The seller's objective is to maximize the present value of the expected profit from the car. In the following, we first introduce the key components of our model before we formally set it up.

Seller learning

We adopt the Bayesian Normal Learning framework to capture the seller's learning process. We assume that $\xi_j \sim N(0, \sigma_\xi^2)$. The seller inspects the car before setting the price for the first buyer. We quantify the result of the initial assessment as a signal, y_0 , drawn from $N(\xi_j, \sigma_0^2)$. With y_0 , the seller updates her belief about ξ_j using the Bayes' rule.

For seller learning in the sale process, we intend to capture two common sources of learning in the market. First, the seller can make inferences about ξ_j by simply observing the event of no sale. The seller would adjust her belief about ξ_j downward every time she sees a buyer walking away from the car. Another possibility is that the seller may get some information about buyers' values for car j while talking to them. Buyers' values contain direct information about ξ_j . To capture the two sources of learning, we assume that, if a buyer decides not to buy a car, he reveals his value for the car to the seller. Then, in effect, the seller receives a signal $y_{ij} \equiv \xi_j + \varepsilon_{ij}$ about ξ_j from every buyer who chooses not to buy car j . Similarly, the seller updates her belief about ξ_j using the Bayes' rule every time she receives a new signal.

The above specification of the learning process covers both sources of learning that we want to capture. Obviously, it covers learning via communications with buyers. In fact, it also covers learning about ξ_j from observing no sale, because, in the inference of ξ_j , $\xi_j + \varepsilon_i$ is a sufficient statistic for the event of no sale. The specification seems reasonable from both the buyers' and the seller's perspectives. Buyers have no disincentive to truthfully reveal their values because they are making one-shot purchase decisions, and the prices they face are nonnegotiable. The seller should also have the incentive to elicit the values from buyers, which improves her information about ξ_j . As we will discuss in more detail below, this specification of the learning process is also important for keeping the model tractable.

For later reference, we use $y^i \equiv (y_0, \dots, y_{i-1})$ to denote the vector of signals that the seller receives before the arrival of buyer i . And let $(\mu(y^i), \sigma_i^2)$ be the mean and variance of the seller's posterior belief about ξ_j after observing y^i . To keep track of the seller's belief, we need to know only $(\mu(y^i), \sigma_i^2)$, since both the prior beliefs and the signals have normal distributions. For simplicity, we sometimes write μ_i in place of $\mu(y^i)$.

Menu cost

The seller faces some costs to change prices. The most direct one is the physical cost of updating the prices posted on cars and in advertisements. However, as every used car is unique, the price-adjustment process has to be customized for each individual car. So, more importantly, the menu cost could represent the cost of assessing new information and coming up with a new price given the updated belief about ξ_j for each individual car. The magnitude of the menu cost can depend on the seller's ability to assess new information and the information-management and pricing system that the seller has in place. We use a scalar parameter ϕ_1 to denote the seller's cost of resetting price once for a car.

Competition and inventory management

Let the number of cars in the current inventory be J_{i1} . These cars compete for the current demand. The competition that a car faces increases with the number of other cars currently available. Naturally, the price that the seller sets for a car would respond to such a competition effect.

For the purpose of managing inventory, the seller also needs take the current inventory into account. When the current inventory is high, the chances of future stock-out is lower. So, the seller may set prices lower to sell faster. But when the current inventory is low, the chance of future stock-out is higher, and, thus, the seller may want to set prices higher to balance the current profit and potential future loss due to stock-out. As a convenient way of capturing the seller's need to manage inventory in our model, we let J_{i1} directly enter the seller static profit function.

In practice, the seller might know in advance that some cars have been acquired and will be added to the inventory. Let J_{i2} denote the number of arriving cars that the seller observes in advance. Observing J_{i2} also affects current prices, because it helps the seller predict future competition and future need to manage inventory. To allow for this possibility, we let J_{i2} enter the model as a state variable, though it does not directly affect the seller's current expected profit. We assume that J_{i2} evolves as a first-order Markov process, and that, conditional on J_{i2} , the current inventory J_{i1} also evolves as a first-order Markov process. We define $J_i \equiv (J_{i1}, J_{i2})$ as the vector of the current inventory and arriving inventory.

Dynamic pricing in the presence of seller learning

We consider two scenarios for the subsequent development of the model. One has the seller learning about ξ_j in the sale process, and the other does not. The comparison allows us to highlight the important properties of pricing behavior in the presence of seller learning and the complications involved in empirically quantifying the value of the information about ξ_j to the seller. In the following, we first introduce our model of dynamic pricing with seller learning. We will use the function $D(p_{ij}, \xi_j, \varepsilon_{ij}) \equiv E_{(\varepsilon_{ia}, \varepsilon_{i0})} I_{ij}(p_{ij}, \xi_j, \varepsilon_{ij}, \varepsilon_{ia}, \varepsilon_0)$ to denote buyer i 's probability of buying car j conditional on $(p_{ij}, \xi_j, \varepsilon_{ij})$, emphasizing its dependence on the price p_{ij} , the car's quality ξ_j , and the buyer's idiosyncratic preference shock ε_{ij} . In what follows, we will suppress the car index j for notational simplicity.

Let $p_i : R^{i+2} \rightarrow R^+$, be a pricing function that maps (y^i, J_i) to a price. Then, the seller's pricing strategy can be expressed as $(p_i)_{i=1}^\infty$, which maps the seller's latest information in every stage to a price. Formally, the seller's profit-maximization problem

can be written as follows:

$$\begin{aligned}
& \max_{(p_i)_{i=1}^{\infty}} E_{\xi} E_{y^1 | \xi} E_{(\varepsilon_i)_{i=1}^{\infty}} E_{(J_i)_{i=2}^{\infty} | J_1} \sum_{i=1}^{\infty} s_i \delta^{i-1} \pi_i(p_i, \xi, \varepsilon_i, J_i) \\
\text{s.t. } \pi_i &= -\phi_1 1\{p_i \neq p_{i-1}\} + (p_i + \phi_2 J_{i1}) D(p_i, \xi, \varepsilon_i) \\
y^{i+1} &= (y^i, y_i) \\
y_i &= \xi + \varepsilon_i
\end{aligned}$$

where s_i indicates the availability of the car at the beginning of period i , and δ is the seller's discount factor.

The above problem is difficult to solve directly. However, given our specification of the learning process, it can be transformed into a sequential optimization problem. First, note that the above optimization problem can be reformulated as follows:

$$\begin{aligned}
& \max_{(p_i)_{i=1}^{\infty}} \left\{ E_{\xi} E_{y^1 | \xi} \left(E_{(\varepsilon_i)_{i=1}^{\infty}} E_{(J_i)_{i=2}^{\infty} | J_1} \sum_{i=1}^{\infty} s_i \delta^{i-1} \pi_i(p_i, \xi, \varepsilon_i, J_i) \right) \right\} \quad (1) \\
& = E_{y^1} \max_{(p_i)_{i=1}^{\infty}} \left\{ E_{\xi | y^1} \left(E_{(\varepsilon_i)_{i=1}^{\infty}} E_{(J_i)_{i=2}^{\infty} | J_1} \sum_{i=1}^{\infty} s_i \delta^{i-1} \pi_i(p_i, \xi, \varepsilon_i, J_i) \right) \right\}
\end{aligned}$$

where the equality follows by changing the order of integration. The equation says that the seller maximizes her expected profit from selling the car if and only if she maximizes her expected profit based on her updated belief about ξ after receiving signal y^1 for every value of y^1 .

Furthermore, given the vector of signals y^k that the seller has received before the arrival of buyer k , we have:

$$\begin{aligned}
& \max_{(p_i)_{i=k}^{\infty}} \left\{ E_{\xi | y^k} \left(E_{(\varepsilon_i)_{i=k}^{\infty}} E_{(J_i)_{i=k+1}^{\infty} | J_k} \sum_{i=k}^{\infty} s_i \delta^{i-k} \pi_i(p_i, \xi, \varepsilon_i, J_i) \right) \right\} \\
& = \max_{(p_k)} \left\{ E_{\xi | y^k} E_{\varepsilon_k} \pi_k(p_k, \xi, \varepsilon_k, J_k) + \max_{(p_i)_{i=k+1}^{\infty}} \right. \\
& \quad \left. \left\{ E_{\xi | y^k} E_{\varepsilon_k} \left((1 - D(p_k, \xi, \varepsilon_k)) \delta E_{(\varepsilon_i)_{i=k+1}^{\infty}} E_{(J_i)_{i=k+1}^{\infty} | J_k} \sum_{i=k+1}^{\infty} s_i \delta^{i-(k+1)} \pi_i(p_i, \xi, \varepsilon_i, J_i) \right) \right\} \right\} \\
& = \max_{(p_k)} \left\{ E_{\xi | y^k} E_{\varepsilon_k} \pi_k(p_k, \xi, \varepsilon_k, J_k) + E_{y^{k+1} | y^k} (1 - D(p_k, \xi, \varepsilon_k)) \right. \\
& \quad \left. \delta \max_{(p_i)_{i=k+1}^{\infty}} E_{\xi | y^{k+1}} \left\{ E_{(\varepsilon_i)_{i=k+1}^{\infty}} E_{(J_i)_{i=k+1}^{\infty} | J_k} \sum_{i=k+1}^{\infty} s_i \delta^{i-(k+1)} \pi_i(p_i, \xi, \varepsilon_i, J_i) \right\} \right\} \quad (2)
\end{aligned}$$

where the second equality follows by noting that $y^{k+1} = (y^k, y_k) = (y^k, \xi + \varepsilon_k)$ and that $D(p_k, \xi, \varepsilon_k)$ does not vary with ξ after conditioning on the signal $y_k = \xi + \varepsilon_k$. In

the expression following the first equality sign above, both the conditional continuing probability, $1 - D(p_k, \xi, \varepsilon_k)$, and the total expected future payoffs depend on (ξ, ε_k) .⁷ By conditioning on y_k , we separate the continuing probability and the continuation value of selling to future buyers.⁸ Taken together, reformulations (1) and (2) imply that the seller's original profit optimization problem can be transformed into a sequential optimization problem, which has a Bellman Equation representation. Let us define the following value function:

$$\tilde{V}^k(\tilde{S}_k) = \max_{(p_i)_{i=k}^{\infty}} E_{\xi|y^k} E_{\varepsilon_k} E_{(\varepsilon_i)_{i=k}^{\infty}} E_{(J_i)_{i=k+1}^{\infty}|J_k} \sum_{i=k}^{\infty} \delta^{i-k} s_i \pi_i(p_i, \xi, \varepsilon_i, J_i)$$

where $\tilde{S}_k \equiv (y^k, J_k, p_{k-1})$. Then, the seller's profit-optimization problem has the following Bellman Equation representation:

$$\begin{aligned} \tilde{V}^k(\tilde{S}_k) &= \max_{p_k} \left\{ E_{\xi|y^k} E_{\varepsilon_k} \pi_k(p_k, \xi, \varepsilon_k, J_k) + E_{y^{k+1}|y^k} (1 - D(p_k, \xi, \varepsilon_k)) \delta E_{J_{k+1}|J_k} \tilde{V}^{k+1}(\tilde{S}_{k+1}) \right\} \\ s.t. \pi_k &= -\phi_1 1\{p_k \neq p_{k-1}\} + (p_k + \phi_2 J_{k1}) D(p_k, \xi, \varepsilon_k) \\ y^{k+1} &= (y^k, y_k) \\ y_k &= \xi + \varepsilon_k \end{aligned}$$

The above representation makes the seller's trade-offs in setting prices transparent. When the seller chooses the price p_k , she expects not only that p_k determines the current expected payoff, but also that, with probability $1 - D(p_k, \xi, \varepsilon_k)$, she will receive the new signal $y_k = \xi + \varepsilon_k$ and continue to sell the car to future buyers. Note that the continuing probability, $1 - D(p_k, \xi, \varepsilon_k)$, is increasing in p_k and decreasing in $\xi + \varepsilon_k$. Suppose that \tilde{p}_k is the price that maximizes the current expected payoff. Then, setting the current price higher than \tilde{p}_k increases the probability of receiving the continuation value. Furthermore, higher current price implies that the distribution of signal y_k , conditional on continuing, would be more concentrated on higher values, and, consequently, the continuation value of selling to future buyers is also higher. Therefore, the seller's trade-off is between maximizing the current expected payoff and "increasing the continuation value and the probability of receiving it."

⁷The future payoffs depend on both ξ and ε_k , because the next price is set after the seller receives the signal $y_k = \xi + \varepsilon_k$ and ξ directly enters the expected future payoffs.

⁸The assumption that buyers reveal their values to the seller, which effectively reveals $\xi + \varepsilon_k$, is partly motivated by the need to transform the original optimization problem into a sequential optimization problem. Had we assumed that the signal that the seller received in the learning process is something other than $\xi + \varepsilon_k$, we would not be able to separate the continuing probability and the continuation value by conditioning on the signal, and the original profit-maximization problem cannot be transformed into a sequential optimization problem either.

Assuming stationarity for the optimal pricing strategy, we have the following slightly more concise Bellman equation representation for the seller's original profit maximization problem:

$$\begin{aligned}\tilde{V}(S_k) &= \max_{p_k} \left\{ E_{\xi|y^k} E_{\varepsilon_k} \pi_k(p_k, \xi, \varepsilon_k, J_k) + E_{y^{k+1}|y^k} (1 - D(p_k, \xi, \varepsilon_k)) \delta E_{J_{k+1}|J_k} \tilde{V}(S_{k+1}) \right\} \\ s.t. \pi_k &= -\phi_1 1\{p_k \neq p_{k-1}\} + (p_k + \phi_2 J_{k1}) D(p_k, \xi, \varepsilon_k) \\ \mu_{k+1} &= \frac{\sigma_k^2 y_k + \sigma_\varepsilon^2 \mu_k}{\sigma_k^2 + \sigma_\varepsilon^2} \\ \sigma_{k+1}^2 &= \frac{\sigma_k^2 \sigma_\varepsilon^2}{\sigma_k^2 + \sigma_\varepsilon^2}\end{aligned}$$

where $S_k \equiv ((\mu(y^k), \sigma_k), J_k, p_{k-1})$, and the value function depends on y^k only through $(\mu(y^k), \sigma_k^2)$, the mean and variance of the seller's current belief about ξ . Given the above representation, the seller's profit-maximization problem is, in essence, a stochastic optimal-stopping-time problem with learning. By setting prices, the seller controls the probability of stopping (i.e., sale) given her latest information about ξ .

Dynamic pricing without learning

In the case without seller learning, the seller still gets an initial signal y_0 about ξ through examining the car, but she does not learn about ξ further in the sale process. Let $p_i : R^3 \rightarrow R^+$ be the seller's pricing function, which maps (y_0, J_i) to a price for buyer i . Similarly, the seller's profit-maximization problem can be written as:

$$\begin{aligned}\max_{(p_i)_{i=1}^\infty} E_\xi E_{y^1|\xi} E_{(\varepsilon_i)_{i=1}^\infty} E_{(J_i)_{i=2}^\infty|J_1} \sum_{i=1}^\infty s_i \delta^{i-1} \pi_i(p_i, \xi, \varepsilon_i, J_i) \\ s.t. \pi_i &= -\phi_1 1\{p_i \neq p_{i-1}\} + (p_i + \phi_2 J_{i1}) D(p_i, \xi, \varepsilon_i)\end{aligned}$$

Let us define the following value function:

$$\hat{V}(J_k) = \max_{(p_i)_{i=k}^\infty} E_{\xi|y^1} E_{(\varepsilon_i)_{i=k}^\infty} E_{(J_i)_{i=k+1}^\infty|J_k} \sum_{i=k}^\infty s_i \delta^{i-k} \pi_i(p_i, \xi, \varepsilon_i, J_i)$$

Then, the Bellman equation for the seller's profit-optimization problem can be written as follows:

$$\begin{aligned}\hat{V}(J_k) &= \max_{p_k} E_{\xi|\mu(y^1), \sigma_1} E_{\varepsilon_k} \left(-\phi_1 1\{p_k \neq p_{k-1}\} + (p_k + \phi_2 J_{k1}) D_k + (1 - D_k) \delta E_{J_{k+1}|J_k} \hat{V}(J_{k+1}) \right) \\ s.t. D_k &= D(p_k, \xi, \varepsilon_k)\end{aligned}$$

Note that the seller's belief about ξ is always conditioned on $(\mu(y^1), \sigma_1)$ because of no learning.

4.4 The Optimal Pricing Strategy

In this subsection, we discuss the properties of the optimal pricing strategy in the presence of seller learning. The price affects the seller’s expected profit in three ways. First, it determines the expected profit of selling to the current buyer. Second, it controls the probability of continuing (i.e., no sale) and getting the continuation value. Third, it affects the gain from learning, conditional on continuing. These roles of price heavily shape the price dynamics under the optimal pricing strategy.

Pricing dynamics in the presence of learning

When setting prices in a dynamic process, the seller faces the trade-off between maximizing the current expected profit and getting higher expected option values. Because the seller has the option of selling to future buyers, the optimal price would always be higher than the price that maximizes the current expected profit. In this context, learning influences price via multiple channels.

As is well understood, learning changes the seller’s belief about ξ , which directly affects the option value and the optimal price. Notice that, in our case, learning happens only when the sale process continues, which is more likely when $y_0 > \xi$ —that is, the seller’s initial assessment overestimated the car’s condition ξ . As a result, more often than not, learning leads to a more pessimistic belief about ξ and decreasing option values. Therefore, learning coupled with selection on ξ and y_0 creates a downward trend in the optimal prices.

Were this the only impact of learning on prices, the magnitude of price changes over time would be informative of the amount of learning in the sale process, and of the quality of the seller’s initial information about ξ . However, the following two effects of learning on price dynamics show that the task of inferring the seller’s initial information about ξ is more complicated. They also highlight the importance of developing a full-fledged dynamic model of pricing in order to quantify the value of the information about item-specific demand.

First, in the presence of learning, the dynamics in the value of new information generates steeper drops in the optimal prices. In general, the value of new information leads to higher option values and higher optimal prices. It is straightforward to show that we have $\hat{V}(\mu(y^1), \sigma_1) < E_{y^2|y^1} \tilde{V}(\mu(y^2), \sigma_2)$,⁹ and $\frac{\partial p_1}{\partial EV} > 0$. That is, given the same initial assessment result y^1 , the optimal initial price is higher in the case with learning than in the case without learning. However, as the seller becomes more informed, the value of new information diminishes over time. Figure 4 plots the sequence of ratios of the standard deviation of the seller’s belief in the current period to that in the following period. It shows that the impact of new information on the seller’s belief

⁹The inequality is also a direct implication of the well-known Blackwell’s Theorem.

drops quickly over time. Overall, the diminishing value of new information leads to decreasing option values and decreasing optimal prices.¹⁰

Second, the impact of current price on the gain from subsequent learning generates additional dynamics in optimal prices. As we noted in the model section, the current price affects the distribution of the new signal that the seller receives, conditional on continuing. This creates an incentive for the seller to set higher prices because a higher current price implies a larger gain from learning and higher option value. Such an incentive is strong initially, and then goes down as learning becomes less important over time. So, this would also lead to a higher optimal initial price and decreasing optimal prices subsequently. Mason and Välimäki (2011) make a similar point in a learning model with binary signals. They show that such an effect of price on learning increases the initial price relative to the case without learning.

One additional point about the learning effects discussed above is worth mentioning. The magnitude of the learning effect depends on the value of new information obtained in the learning process. The seller's expected continuation value in the Bellman Equation is a convex function in its posterior belief about ξ_j (see Mason and Välimäki (2011)). The convexity reflects how much the optimal price changes with the posterior belief about ξ_j . Larger convexity in the expected continuation value function implies larger value for the new information.

Overall patterns of price dynamics

We make three observations about the overall pattern of price dynamics under the optimal dynamic pricing strategy. First, the optimal prices drop, on average, over time. The selection on ξ and y_0 is a major force that drives down the optimal prices. The latter two effects of learning mentioned above lead to higher initial prices and larger subsequent price drops.

Second, an individual sequence of optimal prices can go either up or down as a result of learning and changes in inventory. Because of selection, the seller's belief is more likely to be revised downward in the process. However, as the value of the seller's outside options are also random variables, the seller could have her belief revised upward when the signal about ξ is high, and the outside option value is also high enough so that the current buyer decides not to buy the car under consideration.

Lastly, the optimal prices would be sticky, staying constant for some time, because of the menu cost. Given the diminishing impact of learning, the existence of menu cost implies that price will change more frequently in the first few days because of the larger impact of learning in the early days.

¹⁰In simulated optimal price sequences, we can see cases in which the price significantly drops even though the expected ξ actually increases. The reason, as we argued above, is that the option value drops significantly as the value of new information quickly diminishes.

5 Estimation and Identification

5.1 Estimating the Demand Model

There are two main issues in the estimation of the demand model. One is that the price is potentially endogenous. The seller has some information about ξ_j when she sets price, which makes price potentially correlated with ξ_j . The other is that we need to deal with the unobserved heterogeneity ξ_j . Following the suggestion of Petrin and Train (2010),¹¹ we use the control function approach to deal with the price endogeneity, and capture the unobserved heterogeneity as a random effect. In the following, we describe in detail how we estimate the demand model.

As we do not observe any data on buyers, we make the normalization assumption that one and only one buyer looks at a car each day. So, the number of days that a car has stayed on the market is the same as the number of buyers that have inspected the car before the current buyer. In order to minimize the impact of price changes on demand estimates, we use only the first two days' data of each car for estimating the demand model.¹² In the discussion below, we assume that the price stays constant for the first two days.

Let $X_j \equiv (X_{1j}, X_{2j})$ be the exogenous variables, among which X_{2j} is excluded from buyers' utility functions. The constant term is the first element in X_{1j} .¹³ Suppose that we have the following reduced-form pricing equation for the initial price:

$$p_{1j} = X_j \varphi + \epsilon_j$$

and that $(\xi_j, \epsilon_j) \perp X_j$, and $(\xi_j, \epsilon_j) \sim Normal \left(0, \begin{pmatrix} \sigma_\xi^2 & \sigma_{\xi\epsilon}^2 \\ \sigma_{\xi\epsilon}^2 & \sigma_\epsilon^2 \end{pmatrix} \right)$. Note that $\sigma_{\xi\epsilon}^2 \neq 0$ is the source of endogeneity of price. Then, it follows that $\xi_j = \frac{\sigma_{\xi\epsilon}^2}{\sigma_\epsilon^2} \epsilon_j + \eta_j$, where $\eta_j \perp (X_j, \epsilon_j)$, and we have that $\eta_j \perp (\epsilon_j, p_{1j}, X_j)$ and $\eta_j \sim N \left(0, \sigma_\xi^2 - \frac{\sigma_{\xi\epsilon}^4}{\sigma_\epsilon^2} \right)$. So, we can rewrite buyer i 's value for buying car j as follows:

$$\begin{aligned} v_{ij} &= X_{1j} \beta + p_{1j} \alpha + \xi_j + \varepsilon_{ij} \\ &= X_{1j} \beta + p_{1j} \alpha + \frac{\sigma_{\xi\epsilon}^2}{\sigma_\epsilon^2} \epsilon_j + \eta_j + \varepsilon_{ij} \end{aligned}$$

For the competition effect, we focus on competition within narrowly defined segments—e.g., Japanese midsize cars. Let N_{ij} indicate the number of cars in the segment of car j on day i . We approximate the mean value of buying other cars as a function of the

¹¹See, also, discussions by Villas-Boas (2007).

¹²For the sample used in the demand estimation, around four percent of the cars have their prices changed on the second day.

¹³The coefficient of the constant is β_0 .

number of other cars in car j 's segment and a choice-specific constant:

$$u_{ia} = \beta_0 + \bar{u}_a + \rho \log(N_{ij} - 1)$$

where \bar{u}_a is the choice-specific constant. The approximation may be interpreted as aggregating the choices of buying other cars as a single option and implementing the aggregation similar to McFadden et al. (1978). This simplifying assumption is motivated mainly by the need to keep the state space tractable for estimating the model of dynamic pricing.

We use future additions to the inventory of the same segment as the excluded instrumental variable, X_{2j} , for price. As we have argued above, the arriving inventory would affect the current price given that the seller is forward-looking. However, cars entering the market in the future should not affect the current buyers' purchase decisions since they make a one-shot purchase decision and then exit the market.

Let τ_j be the number of days that car j took to sell, and define $T_j \equiv \min\{2, \tau_j\}$. Then, the data we use in the demand estimation are $\left(X_j, p_j, (I_{ij})_{i=1}^{T_j}\right)_{j=1}^N$, and the set of parameters to be estimated is $\theta_d \equiv (\beta, \alpha, \bar{u}_a, \rho, \sigma)$. For computing the likelihood, let h_{ij} denote the sale probability of car j on i th day:

$$h_{ij} \equiv \Pr(I_{ij} = 1 | X_j, p_{ij}, \epsilon_j, \eta_j; \theta_d)$$

where I_{ij} is the indicator of car j being sold on the i 'th day. Then, we have the following expression for the likelihood of the observation of car j :

$$L\left((I_{ij})_{i=1}^{T_j} | X_j, p_{ij}, \epsilon_j, \eta_j; \theta_d\right) = \prod_{i=1}^{T_j} h_{ij}^{I_{ij}} (1 - h_{ij})^{1-I_{ij}}$$

As we do not observe η_j , we cannot directly compute the above likelihood for car j . But, we observe each car from the beginning, and $\eta_j \perp (\epsilon_j, p_{1j}, X_j)$. So, we can treat η_j as a car random effect, and concentrate out η_j to get the following observable likelihood function for car j :

$$L\left((I_{ij})_{i=1}^{T_j} | X_j, p_{ij}, \epsilon_j; \theta_d\right) = \int L\left((I_{ij})_{i=1}^{T_j} | X_j, p_{ij}, \epsilon_j, \eta_j; \theta_d\right) dP(\eta_j)$$

We can then estimate the structural parameters by using the method of Maximum Likelihood (ML) as follows:

$$\hat{\theta}_d = \arg \max_{\theta_d} \sum_{j=1}^J \log L\left((I_{ij})_{i=1}^{T_j} | X_j, p_{ij}, \epsilon_j; \theta_d\right)$$

5.2 Estimating the Dynamic Pricing Model

The estimation of the dynamic pricing model is a bit more complicated. In general, we use the method of Simulated ML to estimate the structural parameters, $\theta_s \equiv$

$(\sigma_0, \phi_1, \phi_2)$, in the model. We focus on the first T days' data of each car, and define again $T_j \equiv \min\{T, \tau_j\}$. Limiting the number of days used in the estimation does not affect identification, but helps reduce computational burden.¹⁴ Then, we have the following expression for the likelihood of the observed price sequence of car j :

$$\begin{aligned} & l(p^{T_j}, J_j | \theta_s) \\ &= \prod_{i=1}^{T_j} l(p_i | p^{i-1}, J_{ij}; \theta_s) \cdot \Pr(J_{ij} | J_{i-1,j}) \\ &= \prod_{i=1}^{T_j} \int l(p_i | p^{i-1}, y^i, J_{ij}; \theta_s) f(y^i | p^{i-1}) dy^i \cdot \Pr(J_{ij} | J_{i-1,j}) \end{aligned}$$

where $J_j \equiv (J_{ij})_{i=1}^{T_j}$, $p^i \equiv (p_i)_{i=1}^{T_j}$; $f(y^i | p^{i-1})$ is the probability density function of y^i conditional on p^{i-1} ; and $\Pr(J_{ij} | J_{i-1,j})$ is the conditional transition probability of J_{ij} . In the last expression above, $l(p_i | p^{i-1}, y^i, J_{ij}; \theta_s)$ is the conditional likelihood that we can compute using the corresponding optimal pricing strategy. We integrate y^i out of $l(p_i | p^{i-1}, y^i, J_{ij}; \theta_s)$ to get the observable likelihood $l(p_i | p^{i-1}, J_{ij}; \theta_s)$. Note that the ML estimator of θ_s does not depend on $\Pr(J_{ij} | J_{i-1,j})$ directly.¹⁵ So, we can estimate θ_s by using the method of ML as follows:

$$\hat{\theta}_s = \arg \max_{\theta_s} \sum_{j=1}^N \sum_{i=1}^{T_j} \log \left(\int l(p_i | p^{i-1}, y^i, J_{ij}; \theta_s) f(y^i | p^{i-1}) dy^i \right) \quad (3)$$

Following Rust (1987), we compute the above ML estimator by using the Nested Fixed Point algorithm. The algorithm involves an inner loop and an outer loop. The inner loop solves the dynamic pricing model for any given θ_s , and the outer loop searches over the space of θ_s to look for the $\hat{\theta}_s$ that maximizes the likelihood of the data.

We use the Parametric Policy Iteration method (c.f. Benítez-Silva, Hall, Hitsch, Pauletto, Brook, and Rust (2000)) to solve the dynamic pricing model numerically, where we parameterize the value function using the Chebyshev polynomials and iterate over the policy function until convergence. The details of the numerical solution method is deferred to the Appendix. Let us write the optimal pricing strategy for a given θ_s as $\wp((\mu(y^i), \sigma_i^2), J_{ij}, p_{i-1,j}; \theta_s)$, which we abbreviate as $\wp(y^i, p_{i-1,j})$ for the simplicity of notation.

Computing the log-likelihood function in (3) is difficult, because it involves high dimensional integrations over the conditional distributions of serially correlated signals.¹⁶

¹⁴In computing the likelihood of the observation of a car, we need to compute the model-predicted optimal price for each day observed for the car. The computation is costly, especially because the model-predicted optimal price for each day is different for different cars. As we argued in the model section, this is also partly why the cost of adjusting prices may be nontrivial for the seller.

¹⁵ $\Pr(J_{ij} | J_{i-1,j})$ is directly estimated from data. It is used only when we numerically solve the dynamic pricing model.

¹⁶Without conditioning on ξ_j , the signals y_i are correlated across periods.

As there is no analytical expression for the integration, we compute it via simulation. In particular, we use the method of Sampling and Importance Re-sampling (SIR) to simulate the integrations.^{17 18} In the following, we describe in detail how we simulate $l(p_i|p^{i-1}, J_{ij}; \theta_s)$ by using the method of SIR. The dependence of the likelihoods on θ_s and J_{ij} will be suppressed for the simplicity of notation.

The likelihood for the first period is straightforward to simulate. For $l(p_1)$, we have:

$$\begin{aligned} l(p_1) &= \int f(p_1|y_0) f(y_0) dy_0 \\ &= \int 1\{\wp(y_0) = p_1\} f(y_0) dy_0 \\ f(y_0) &= \int f(y_0|\xi_j) f(\xi_j) d\xi_j \\ &= \frac{1}{\sqrt{\sigma_\xi^2 + \sigma_0^2}} \phi\left(\frac{y_0}{\sqrt{\sigma_\xi^2 + \sigma_0^2}}\right) \end{aligned}$$

where ϕ is the probability density function of the standard normal distribution. Then, we can draw a random sample of $\{y_{0s}\}_{s=1}^{N_s} \sim f(y_0)$, and simulate $l(p_1)$ as follows:

$$l(p_1) = \frac{1}{N_s} \sum_{s=1}^{N_s} 1\{\wp(y_{0s}) = p_1\}$$

In the following, we show how to simulate each of the conditional likelihoods. For $l(p_2|p_1)$, we have:

$$\begin{aligned} l(p_2|p_1) &= \int l(p_2|y^2, p_1) f(y^2|p_1) dy^2 \\ f(y^2|p_1) &= f((y_1, y_0)|p_1) \\ &= f(y_1|y_0, p_1) f(y_0|p_1) \\ &= f(y_1|y_0) f(y_0|p_1) \end{aligned}$$

¹⁷See Rubin (1988). See also Fernandez-Villaverde and Rubio-Ramirez (2007), Flury and Shephard (2008) and Gallant, Hong, and Khwaja (2009) for examples of applying SIR to simulate the likelihood function when estimating dynamic models.

¹⁸Alternatively, we can express $l(p^{T_j}, J_j|\theta_s)$ as follows:

$$l(p^{T_j}, J_j|\theta_s) = \int l(p^{T_j}, J_j|y^{T_j}, \theta_s) f(y^{T_j}) dy^{T_j}$$

where $f(y^{T_j})$ is the probability density function of y^{T_j} . Then we may compute $l(p^{T_j}, J_j|\theta_s)$ by simulation, using random draws of y^{T_j} directly from its distribution. However, given the high dimensionality of the integration, it takes a very large number of random draws to simulate the integration with reasonable precision. Furthermore, the method becomes almost infeasible especially because we have to compute the optimal prices for each given random draw of y^{T_j} for every car.

where the last equality follows by noting that p_1 is independent of y_1 after conditioning on y_0 ,¹⁹ and

$$\begin{aligned} f(y_1|y_0) &= \int f(y_1|\xi_j) f(\xi_j|y_0) d\xi_j \\ &= \frac{1}{\sqrt{\sigma_\varepsilon^2 + \sigma_1^2}} \phi\left(\frac{y_1 - \mu_1(y_0)}{\sqrt{\sigma_\varepsilon^2 + \sigma_1^2}}\right) \\ f(y_0|p_1) &= \frac{f(p_1|y_0) f(y_0)}{f(p_1)} \\ &\propto f(p_1|y_0) f(y_0) \end{aligned}$$

Now, to simulate $l(p_2|p_1)$, we need to draw a random sample of $\{y_s^{2|1}\}_s \sim f(y^2|p_1)$, which involves a filtering step and a prediction step. In the filtering step, we draw a random sample $\{\tilde{y}_{0s}\}$ of size N_s via importance re-sampling from $\{y_{0s}\}_{s=1}^{N_s} \sim f(y_0)$, using $f(p_1|y_{0s}) = 1\{\varphi(y_{0s}) = p_1\}$ as the importance sampling weight. Then, for each given \tilde{y}_{0s} , we draw one y_{1s} from the distribution of $f(y_1|y_0)$, which is called the ‘‘prediction’’ step. Then, $\{y_s^{2|1}\}_{s=1}^{N_s} \equiv \{(y_{1s}, \tilde{y}_{0s})\}_{s=1}^{N_s}$ is a random sample from $f(y^2|p_1)$, and we can simulate $l(p_2|p_1)$ as follows:

$$l(p_2|p_1) = \frac{1}{N_s} \sum_{s=1}^{N_s} 1\{\varphi(y_s^{2|1}, p_1) = p_2\}$$

It is important to note that y_i is serially correlated across periods and we have $f(y_1|y_0) \neq f(y_1)$.

Recall that $y^i \equiv (y_0, \dots, y_{i-1})$ is the vector of signals that the seller has received before the arrival of buyer i . Suppose that $\{y_s^{i-1|i-2}\}_{s=1}^{N_s}$ is a random sample we have drawn from the distribution of $f(y^{i-1}|p^{i-2})$. In general, for the conditional likelihood $l(p_i|p^{i-1})$, we have the following relations:

$$\begin{aligned} l(p_i|p^{i-1}) &= \int l(p_i|y^i, p^{i-1}) f(y^i|p^{i-1}) dy^i \\ f(y^i|p^{i-1}) &= f((y_{i-1}, y^{i-1})|p^{i-1}) \\ &= f(y_{i-1}|y^{i-1}) f(y^{i-1}|p^{i-1}) \end{aligned}$$

¹⁹Recall that the seller observes only y_0 when she sets price p_1 .

where

$$\begin{aligned}
f(y_{i-1}|y^{i-1}) &= \int f(y_{i-1}|\xi_j) f(\xi_j|y^{i-1}) d\xi_j \\
&= \frac{1}{\sqrt{\sigma_\varepsilon^2 + \sigma_{i-1}^2}} \phi\left(\frac{y_{i-1} - \mu(y^{i-1})}{\sqrt{\sigma_\varepsilon^2 + \sigma_{i-1}^2}}\right) \\
f(y^{i-1}|p^{i-1}) &= \frac{f(p_{i-1}|y^{i-1}, p^{i-2}) f(y^{i-1}|p^{i-2})}{f(p_{i-1}|p^{i-2})} \\
&\propto f(p_{i-1}|y^{i-1}, p^{i-2}) f(y^{i-1}|p^{i-2})
\end{aligned}$$

We can similarly draw $\{y_s^{i|i-1}\}_{s=1}^{N_s}$ based on the random sample of $\{y_s^{i-1|i-2}\}_{s=1}^{N_s}$ through the same two steps, and simulate $l(p_i|p^{i-1})$ as follows:

$$l(p_i|p^{i-1}) = \frac{1}{N_s} \sum_{s=1}^{N_s} 1\left\{\wp(y_s^{i|i-1}, p^{i-1}) = p_i\right\}$$

In our implementation of the above simulation method, we discretize the price using a fine grid.²⁰ Putting all the simulated conditional likelihoods together, we can estimate θ_s using the method of Simulated ML as in (3).

5.3 Identification

In the following, we provide a brief discussion of the identification of the structural parameters in the dynamic pricing model. We do not estimate the daily discount factor δ , but calibrate it to match an annual interest rate of 10%. The variance of ξ , σ_ξ^2 , can be calculated based the parameter estimates of the demand model.

The parameters that we need to estimate in the dynamic pricing model are σ_0 , ϕ_1 and ϕ_2 . The variance of the seller's initial examination signal, σ_0^2 , affects the conditional distribution of the initial price as well as the extent of learning. Therefore, both the conditional variance of the initial price and the total price changes will help identify σ_0^2 . As for the menu cost, ϕ_1 , it determines how frequently the seller changes prices. Thus, the stickiness in prices helps identify ϕ_1 . Lastly, parameter ϕ_2 is identified by the variation in J_{i1} .

6 Empirical Results

6.1 Sample for estimation

In the estimation, we use the 2011 data of the three most popular models at CarMax in all four zip codes. The three models are Honda Accord, Nissan Altima, and Toyota

²⁰We divide the price range in our data into 300 equal intervals and treat the prices within each interval as the same.

Camry. For the demand estimation, we use the data of the second and the third day of each car’s listing period (we treat these two days as the first two days for each car, because no cars are sold on the first listing day in our data). There are 824 unique cars and 1,585 observations. Table 7 describes these cars by model, zip code, and listing day. Out of the 824 cars, 7.52% are sold on the second day, and 4.73% of the remaining 761 cars are sold on the third day (see Table 8). For estimating the dynamic pricing model, we use the data from the second to the 12th day. The 11 days’ data are enough to identify the structural parameters in the dynamic pricing model.

Table 9 summarizes variables for cars used in the estimation. The list price of a car on day 2 ranges from \$7,100 to \$27,000, with the average being \$17,600. The mileage of a car ranges from 2,000 miles to 118,000 miles, and the average is 41,000 miles. 88.34% of cars are between one and six years old.²¹

In the estimation, we define the car segment as cars of the same model. We also experiment with defining the segment more broadly to include all three models (midsize Japanese cars), which produces very similar estimates. For a particular car, we define a number of variables for the current and arriving inventories of cars of the same model. The current inventory is the number of cars of the same model (excluding the current car) available on the same day at the same store; and it ranges from 0 to 19. We also define four future inventory variables: Inventory in one week is defined as the number of cars of the same model arriving at the same store between one and seven days into the future. Inventories in two, three, and four weeks are similarly defined for cars arriving in the second, third and fourth seven-day period in the future.

6.2 Model Estimates

In the following, we present the estimates of the parameters from the demand model and the dynamic pricing model. As we pool the data for three car models in a year and in four different areas in our estimation, we include the brand, time and region dummies in the model to control for the variation in price/demand across model, time and area.

Demand model

Table 10 presents the estimated parameters in the reduced-form pricing equation. As expected, the inventory to be added in the near future has a negative impact on the current price. The impact is most significant for the inventory to be added within a week. The negative impact is consistent with the competition effect and/or the incentive for the seller to avoid stock-out. The estimated coefficients of other variables also seem

²¹CarMax says in its 2011 annual report: “Our primary focus is vehicles that are 1 to 6 years old, have fewer than 60,000 miles and generally range in price from \$11,000 to \$32,000.”

reasonable. For example, older models are listed at lower prices even after controlling for mileage and other attributes. Larger cars and cars with more powerful engines are listed at higher prices.

It is worth noting that the adjusted R-squared of the regression is around 0.8, meaning that there is still a fair amount of variation in price that cannot be explained by the variation in the observable attributes. This is consistent with the fact that car-specific condition is an important determinant of used-car prices.

Table 11 presents the parameters of the demand model, estimated using the control function approach. As mentioned above, the main issue in estimating the demand model is the endogeneity in price. Ignoring the issue could lead to biased estimates of the price coefficient, which could turn out to be statistically insignificant or even positive. Our estimated price coefficient is -4.799 and is statistically significant at the one-percent level. Other estimated coefficients of car attributes in the demand model also seem reasonable. For example, the coefficient of mileage is -1.858 , significant at the ten-percent level; and the coefficient of the indicator of longer wheelbase is 1.500 , significant at the five-percent level.

The estimated standard deviation of the random effect, η , is about 1.79, significant at the one-percent level. The estimated coefficient of the pricing-equation residual is 0.479, though, which is not significant. Combined together, these two parts show the importance of the car-specific heterogeneity in the data.

For parameters associated with the choice of buying other cars of the same segment, the coefficient of $\log(N_{ij} - 1)$ is 1.136 and is significant at one-percent level. It shows that other cars available in the inventory impose a significant competition effect.

Dynamic pricing model

We have three key structural parameters in our dynamic pricing model: the standard deviation of the distribution of the initial-inspection signal (σ_0); the menu cost (ϕ_1); and the inventory effect (ϕ_2). Table 12 presents the estimated parameters. All three parameters are estimated accurately. The estimated σ_0 is 0.301. This means that CarMax's initial examination produces quite precise information about car-specific heterogeneity. This is consistent with the fact that CarMax inspects all of its cars very carefully and uses its assessment result to guide its pricing.

The estimated menu cost is \$315. The menu cost seems much larger than the physical cost of changing prices. We think that it mostly captures the cost of assessing new information and computing a new optimal price based on the seller's updated belief.

The parameter that captures the seller's need to manage inventory, ϕ_2 , is estimated to be -0.0287 . This means that the seller has an incentive to sell cars at lower prices when the inventory is high, and vice versa. This is consistent with the seller's inventory-management goal of avoiding stock-outs.

Overall, our estimates seem to have good face validity. In the following, we use the point estimates of model parameters to conduct our policy experiments.

6.3 Policy Experiments

With the estimates of the structural model, we are ready to quantify the value of the seller’s initial information about car-specific demand and the value of reducing the cost of adjusting prices. In the policy experiments, we simulate the expected profit in the four cases defined by whether initial inspection and/or learning are part of the sale process. Table 13 shows the expected profits in the four cases relative to the fourth case (no inspection and no learning).²² We first focus on the seller’s expected profit in the following three cases: “inspection and learning,” “no inspection but learning” and “no inspection and no learning.” In the latter two cases, the seller’s initial information about ξ is just the population distribution of ξ . We define the value of information as the difference in the expected profits for the first two cases, and the value of learning as the difference in the expected profits for the latter two cases. Our estimated value of information to CarMax is around \$243 per car, and the estimated value of learning to CarMax is around \$109 per car. The backdrop of these values is the average net profit of around \$740 for the cars in our estimation data. Therefore, learning in the sale process does seem to help reduce the negative impact of the initial information limitation for sellers who are not inspecting their cars carefully. However, even with learning, the value of information to CarMax is still quite large. Based on these estimates, the resources spent by CarMax on assessing car-specific demand does seem to have paid off very well, and it may be also worthwhile for other used car dealers in the market to invest in acquiring such information.

We also assess the benefit of streamlining the price-adjustment process to lower the menu cost. As mentioned before, we think of the menu cost mostly as the cost to the seller to evaluate the newly received information and come up with an updated price. Having in place the necessary information infrastructure (which may include the historical sales and demand data and the managerial software and models that computes the optimal prices based on the latest information.) and a streamlined process of collecting and quantifying car and demand information could mean a lower menu cost.

Table 14 shows the simulated expected profits with the estimated parameters and three levels of menu cost. In comparison to the case with the estimated menu cost of \$315, the expected profit increases by around \$103 per car if the menu cost is reduced by half; the expected profit drops by about \$50 if the menu cost is doubled. The improvement in expected profit after reducing the estimated menu cost by half seems

²²In the following, we will refer to the initial inspection simply as inspection, learning in the sale process as learning, the value of the initial information about item-specific demand as the value of information, and the value of learning in the sale process as the value of learning.

quite impressive. It may be worthwhile for CarMax to investigate the possibility of lowering the menu cost so that it can adjust prices more smoothly over time.

7 Conclusion

In the above, we estimated the values of the information about item-specific demand and of lowering menu cost in the setting of dynamic pricing in the used-car market. For this purpose, we developed a structural model of dynamic pricing with seller learning. The model is straightforward to solve and is able to generate the main patterns of price dynamics in the data. With some adjustments, the model could also be applied to price-setting problems in other secondary markets.

Our results show that the potential return to investing in the tools for assessing individual cars and analyzing local demand is large. Given the importance of the market, we hope that our results will draw more attention to the benefits of more-informed pricing in this market.

Appendix

Numerical Solution Method for the Model of Dynamic Pricing with Seller Learning

In this, we describe in detail the numerical method we use to solve the model of dynamic pricing with seller learning presented in the model section. Our objective is to solve the following Bellman equation:

$$\begin{aligned}\tilde{V}(S_k) &= \max_p E_{\xi|y} E_{\varepsilon_k} \left(\phi_1 1\{p_k \neq p_{k-1}\} + (p_k + \phi_2 J_{k1}) D_k + (1 - D_k) \delta E_{y^{k+1}|y^k} E_{J_{k+1}|J_k} \tilde{V}(S_{k+1}) \right) \\ s.t. D_k &= D(p_k, \xi, \varepsilon_k) \\ \mu_{i+1} &= \frac{\sigma_i^2 y_i + \sigma_y^2 \mu_i}{\sigma_i^2 + \sigma_y^2} \\ \sigma_{i+1}^2 &= \frac{\sigma_i^2 \sigma_y^2}{\sigma_i^2 + \sigma_y^2}\end{aligned}$$

where $S_k \equiv ((\mu(y^k), \sigma_k), p_{k-1}, J_k)$. Among the state variables, $(\mu(y^k), \sigma_k, p_{k-1})$ are continuous variables, and J_k is a vector of two discrete state variables.

We use the Parametric Policy Iteration algorithm to solve the Bellman equation. We parameterize the value function by approximating it using a linear combination of continuous basis functions. More specifically, we approximate the value function as follows:

$$\begin{aligned}\tilde{V}(\mu_k, \sigma_k, p_{k-1}, J_k) \\ = \sum_{l=1}^L \psi_l \rho_l(\mu_k, \sigma_k, p_{k-1}, J_k)\end{aligned}$$

where $\rho_l(\mu_k, \sigma_k, p_{k-1}, J_k)$ are multivariate basis functions. In the paper, we use the Chebyshev polynomials as the basis functions.

The policy iteration algorithm involves two steps for each iteration. The first one is the policy-evaluation step for a given strategy $p_k(S)$, which involves solving the following linear functional equation:

$$\tilde{V}(S_k) = E_{\xi|y^k} E_{\varepsilon_k} \left[\phi_1 1\{p_k \neq p_{k-1}\} + (p_k + \phi_2 J_{k1}) D_k + (1 - D_k) \delta E_{y^{k+1}|y^k} E_{J_{k+1}|J_k} \tilde{V}(S_{k+1}) \right]$$

which becomes the following system of equations after substituting in the approximating polynomial:

$$\rho(S_k) \psi = E_{\xi|y^k} E_{\varepsilon_k} \left(\phi_1 1\{p_k \neq p_{k-1}\} + (p_k + \phi_2 J_{k1}) D_k + E_{y^{k+1}|y^k} (1 - D_k) \delta E_{J_{k+1}|J_k} \rho(S_{k+1}) \right) \psi$$

After combining terms with the same coefficients, we have:

$$\begin{aligned}(\rho(S_k) - E_{y^{k+1}|y^k} (1 - D_k) \delta E_{J_{k+1}|J_k} \rho(S_{k+1})) \psi \\ = E_{\xi|y^k} E_{\varepsilon_k} \left(\phi_1 1\{p_k \neq p_{k-1}\} + (p_k + \phi_2 J_{k1}) D_k \right)\end{aligned}$$

Let us use the following notations:

$$\begin{aligned}\varrho(S_k) &\equiv (\rho(S_k) - E_{y^{k+1}|y^k} (1 - D_k) \delta E_{J_{k+1}|J_k} \rho(S_{k+1})) \\ U(S_k) &\equiv E_{\xi|y^k} E_{\varepsilon_k} (\phi_1 1 \{p_k \neq p_{k-1}\} + (p_k + \phi_2 J_{k1}) D_k)\end{aligned}$$

If we pick a grid of N different points of S_k : (S_{k1}, \dots, S_{kN}) , then, we can solve for ψ by using the least squares criterion as follows:

$$\psi^* = \arg \min_{\psi} \sum_{i=1}^N (\varrho(S_{ki}) \psi - U(S_{ki}))^2$$

Given $\tilde{V}(S_k) = \rho(S_k) \psi^*$, we can carry out the policy function improvement step by solving the optimal prices at the grid points as follows:

$$\begin{aligned}p^*(S_{ki}) &= \arg \max_{p_k} E_{\xi|y^k} E_{\varepsilon_k} (\phi_1 1 \{p_k \neq p_{k-1}\} + (p_k + \phi_2 J_{k1}) D_k + \\ &\quad (1 - D_k) \delta E_{y^{k+1}|y^k} E_{J_{k+1}|J_k} (S_{k+1}) \rho(S_k) \psi^*)\end{aligned}$$

We iterate over the above two steps until a convergence criterion for the value function or policy function is satisfied. The fixed point gives the solution to the Bellman equation. The optimal pricing policy can be easily computed with the solution of the value function.

In the above solution method, we need to compute $E_{y^{k+1}|y^k} (1 - D_k)$, and we do not have an analytical formula for it, given our multinomial Probit model. So, we simulate it by using the following formula:

$$\begin{aligned}& E_{y^{k+1}|y^k} D(p_k, \xi, \varepsilon_k) \\ &= \sum_{i=1}^{ns} \Pr(X_1 \beta + \alpha p_k + \xi_i + \varepsilon_{ki} > u_{ka} + \varepsilon_{ka} \ \& \ X_1 \beta + \alpha p_k + \xi_i + \varepsilon_{ki} > \varepsilon_{k0}) \\ &= \sum_{i=1}^{ns} \Pr(X_1 \beta + \alpha p_k + \xi_i + \varepsilon_{ki} > u_{ka} + \varepsilon_{ka}) \Pr(X_1 \beta + \alpha p_k + \xi_i + \varepsilon_{ki} > \varepsilon_{k0}) \\ &= \sum_{i=1}^{ns} \Phi(X_1 \beta + \alpha p_k + \xi_i + \varepsilon_{ki} - u_{ka}) \Phi(X_1 \beta + \alpha p_k + \xi_i + \varepsilon_{ki})\end{aligned}$$

The advantage of the above way of simulating the conditional expected purchase probability is that it ensures that the Bellman equation is smooth and has analytical Jacobian and Hessian. The smoothness property is important for using standard algorithms to search for the optimal prices.

Tables and Figures

Figure 1: An Example of Price Variation by Car Condition from Kelley Blue Book

Home > Car Values > Honda > Accord > 2007 > Style > Options > Condition > LX Sedan 4D

Honda Accord 2007 go

2007 Honda Accord

2006 2007 2008 Style: LX Sedan 4D

edit options | change style

View all 3 photos | Mileage: 68,500 change | Like this car

pricing photos specs kbb expert review consumer reviews ratings compare

Used Car Prices

See Trade-In/Sell Values

Certified Pre-Owned | Suggested Retail | **Private Party**

Excellent \$12,550 Verify Condition	Find Deals Near You see local deals Free VIN Check <input type="text"/> Enter VIN (optional) go <small>No VIN? No Problem!</small> Know Your Credit Score? get yours now <small>freecreditreport.com</small>
Very Good \$12,100	
Good \$11,700	
Fair \$10,400	

Used Cars for Sale near Rockville

PROGRESSIVE


\$99.00

The Name Your Price® Tool. Only from Progressive.

Enter ZIP Code: [Get Your Free Quote](#)


POHANKA HONDA
King of the Bellway

explore




Note: The prices on the left hand side are the KBB “Private Party” prices by car conditions for the 2007 Honda Accord LX sedan with 68,500 miles.

Figure 2: Frequency of price adjustments by time-on-market (TOM): the entire sample



Note: For each value of TOM, the dot is the percentage of remaining cars that adjust prices, and the number that accompanies the dot is the number of cars that remain on the market. The solid lines are the 95% confidence intervals.

Figure 3: Frequency of price adjustments by time-on-the-market (TOM): CarMax



Note: For each value of TOM, the dot is the percentage of remaining cars that adjusted prices, and the number that accompanies the dot is the number of cars that remain on the market on the day. The solid lines are the 95% confidence interval of the percentages.

Figure 4: The Ratio of the Standard Deviation of the Seller's Prior Belief to that of the Posterior Belief




Table 1: Distribution of Time-to-Sell (TOS)

Variable	N	Mean	Sd	Min	p10	p25	p50	p75	p90	Max
Entire sample	13,024	22.93	27.17	1	3	6	13	29	54	274
CarMax	7,630	14.02	13.78	1	2	4	9	19	31	120

Note: TOS of a car is defined as the number of days that the seller took to sell the car.

Table 2: Total Number of Price Changes

(a) Entire Sample

Number of price changes	Freq.	Percent
0	8,276	63.54
1	2,870	22.04
2	1,004	7.71
3	380	2.92
4	228	1.75
5	116	0.89
6	64	0.49
7	45	0.35
8	23	0.18
9	11	0.08
10	3	0.02
11	4	0.03
Total	13,024	100

(b) CarMax

Number of price changes	Freq.	Percent
0	5,350	70.12
1	1,754	22.99
2	402	5.27
3	96	1.26
4	18	0.24
5	7	0.09
6	1	0.01
7	2	0.03
Total	7,630	100

Table 3: Distribution of Price Changes Conditional on Change

(a) One-time Price Changes

Variable	N	Mean	Sd	Min	p25	p50	p75	Max
All price changes	7,815	-508.05	759.44	-10912	-1000	-500	-100	4701
Price increases	874	923.84	667.20	1	500	900	1000	4701
Price decreases	6,941	-688.35	550.15	-10912	-1000	-544	-399	-1

(b) Total Price Changes

Variable	N	Mean	Sd	Min	p25	p50	p75	Max
All price changes	4,655	-914.71	1317.41	-16912	-1396	-758	-146	12210
Price increases	455	1052.50	1107.01	1	500	898	1000	12210
Price decreases	4,200	-1127.83	1151.68	-16912	-1500	-1000	-399	-1

Note: a) One-time price change is defined as (price on day t - price on day $t - 1$); total price change is defined as (price on the last day - price on the first day). b) For one-time changes, we do not include price increases or decreases that are by more than 20%. Most of these cases are entry errors; they are quickly changed back (usually on the next day) within the normal price ranges. c) For total price changes, we do not include price increases or decreases that are by more than 50%. Again, these dropped price changes are mostly entry errors. There are only eight cars dropped.

Table 4: Daily Change of Inventories

Variable	N	Mean	Sd	p25	p50	p75
All dealers	2,183	0.0027	0.0837	-0.0235	0.0000	0.0250
CarMax	364	0.0021	0.0673	-0.0394	0.0076	0.0399
Top six brands of CarMax	364	0.0050	0.1028	-0.0571	0.0000	0.0645

Note: Daily change of inventories is defined as (inventory on day t - inventory on day $t - 1$) / inventory on day $t - 1$ for a dealer.

Table 5: Time-to-sell and Price Adjustments by Dealer

(a) Time-to-Sell (TOS)

Variable	N	Mean	Sd	Min	p25	p50	p75	Max
CarMax	7630	14.02	13.78	1	4	9	19	120
Cook Volkswagen	697	45.95	43.34	1	15	32	64	274
Heritage MileOne Bel Air	958	47.76	45.08	1	13	32	73	272
Jerry's Chevrolet	1193	34.06	33.30	1	10	21	46	271
Jones Junction Toyota Scion	1285	33.75	31.97	1	10	22	47	174
Koons Chevrolet	1261	23.71	18.14	1	9	18	36	85

(b) Total Price Changes

Variable	N	Mean	Sd	Min	p25	p50	p75	Max
CarMax	7630	-153.56	427.30	-6000	-99	0	0	2000
Cook Volkswagen	690	-989.65	1789.1	-16865	-1983	-785	0	8010
Heritage MileOne Bel Air	957	-463.31	1235.95	-16912	-753	0	0	12210
Jerry's Chevrolet	1193	-164.62	563.30	-4197	0	0	0	6510
Jones Junction Toyota Scion	1285	-390.25	1020.36	-10019	-500	0	0	3143
Koons Chevrolet	1261	-1000.95	1449.10	-14000	-2000	-600	0	6000

Table 6: Time-to-sell and Price Adjustments by Brand for CarMax

(a) Time-to-Sell (TOS)

Variable	N	Mean	Sd	Min	p25	p50	p75	Max
Honda Accord	206	13.65	13.82	2	4	9	17	88
Nissan Altima	192	14.70	13.73	1	5	10	21	84
Toyota Camry	135	15.00	12.27	2	5	12	22	66
Honda Civic	163	16.71	16.74	1	4	11	22	92
Chevrolet Impala	116	13.78	12.13	1	5	10	19	81
Chrysler Town & Country	106	12.85	12.05	1	4	9.5	19	58

(b) Total Price Changes

Variable	N	Mean	Sd	Min	p25	p50	p75	Max
Honda Accord	206	-141.57	300.18	-1099	-99	0	0	1000
Nissan Altima	192	-174.84	358.72	-2000	-99	0	0	601
Toyota Camry	135	-137.70	325.28	-1399	0	0	0	601
Honda Civic	163	-101.75	321.94	-2000	0	0	0	1000
Chevrolet Impala	116	-86.97	328.51	-1399	0	0	0	2000
Chrysler Town & Country	106	-89.58	256.72	-1000	0	0	0	601

Table 7: Cars Used in the Estimation

Brand	Zip code				Total
	21162	22911	64055	66202	
<u>Day 2</u>					
Honda Accord	181	38	5	62	286
Nissan Altima	164	35	9	112	320
Toyota Camry	121	29	9	59	218
Total	466	102	23	233	824
<u>Day 3</u>					
Honda Accord	159	38	5	58	260
Nissan Altima	148	35	9	106	298
Toyota Camry	110	27	9	57	203
Total	417	100	23	221	761

Note: The final sample used in the estimation includes 824 cars and 1,585 observations. We use three models listed by CarMax between January and December 2011 in four zip codes. Only the second and the third day of a car's listing period are used for the demand estimation.

Table 8: Car Sales by Day

	<u>Day 2</u>		<u>Day 3</u>	
	Freq.	Percent	Freq.	Percent
Sold	62	7.52%	36	4.73%
Not sold	762	92.48%	725	95.27%

Note: There are 824 cars on the market on day 2, and 761 cars on the market on day 3.

Table 9: Summary Statistics for Cars Used in the Estimation

Variable	Obs	Mean	Std. Dev.	Min	Max
Price (\$10,000)	824	1.76	0.31	0.71	2.70
Mileage (100,000 miles)	824	0.41	0.23	0.02	1.18
Model Year	824	2007.69	1.92	2000	2011
Engine Volume (10 litres)	824	0.26	0.04	0.22	0.35
Wheel Base	824	108.63	1.41	103	110
Sedan	824	0.89	0.32	0	1
Day 2 Day 3					
Current Inventory	1585	7.15	3.91	0	19
Inventory in One Week	1585	1.82	1.68	0	8
Inventory in Two Weeks	1585	3.31	2.46	0	14
Inventory in Three Weeks	1585	1.83	1.75	0	9
Inventory in Four Weeks	1585	2.55	2.18	0	11
Day 2 Day 12					
Price	6,815	1.79	0.29	0.71	2.70
Current Inventory	6,815	7.13	4.04	0.00	19.00

Note: The final sample used in the estimation includes three models (Honda Accord, Nissan Altima, and Toyota Camry) listed by CarMax in 2011 across four zip codes. There are 824 cars, 1,585 observations for the first two days, and 6,815 for the first 11 days.

Table 10: Reduced-form Pricing Equation with Future Inventories as Instruments

Variable	Coefficients	Std
Constant	1.033***	-0.0425
Current Inventory	-0.00242*	-0.00118
Inventory in One Week	-0.00664**	-0.00243
Inventory in Two Weeks	0.000135	-0.00176
Inventory in Three Weeks	0.00138	-0.00224
Inventory in Four Weeks	-0.00376*	-0.00191
Mileage (100,000 miles)	-0.596***	-0.0199
Engine Volume (10 litres)	2.185***	-0.091
Wheelbase = 107	0.168***	-0.0259
Wheelbase = 108-110	0.0534**	-0.0206
Sedan	-0.115***	-0.0177
Transmission	-0.0995***	-0.0203
Model Year = 2003-2004	0.117***	-0.0275
Model Year = 2005-2006	0.203***	-0.0265
Model Year = 2007	0.377***	-0.0283
Model Year = 2008-2009	0.494***	-0.0282
Model Year = 2010	0.603***	-0.0309
Model Year = 2011	0.736***	-0.0424
Honda Accord	0.0830***	-0.0092
Toyota Camry	-0.0626***	-0.0111
Zip Code = 22911	-0.0106	-0.0153
Zip Code = 64055	-0.111***	-0.0247
Zip Code = 66202	-0.0377***	-0.0097
April-May 2011	0.0716***	-0.0107
June-Aug. 2011	0.110***	-0.00994
Sept- Oct 2011	0.0562***	-0.0117
Nov-Dec 2011	-0.0328*	-0.0136
N	1585	
adj. R ²	0.803	

Note: 1) The dependent variable is Price measured in \$10,000. The estimation uses three Japanese models (Honda Accord, Toyota Camry and Nissan Altima) listed by CarMax in four zip codes (21162, 22911, 64055, and 66202) from January to December 2011. The second and third day of a car's listing period are used. 2) Base of Wheelbase is 103-105; base of Model year is 2000-2002; base of Time is Jan.-March 2011. 3) * $p < 0.05$, ** $p < 0.01$, *** $p < 0.1$

Table 11: Multinomial Logit Demand Estimation

Variables	Coefficients	Std
Constant	7.044***	2.434
Price (\$10,000)	-4.799***	1.331
Mileage (100,000 miles)	-1.858*	0.993
Engine Volume (10 litres)	-0.234	2.608
Wheelbase = 107 (base: 103-105)	0.690	0.600
Wheelbase = 108-110	1.500**	0.624
Sedan	-2.105***	0.708
Transmission	0.965	0.836
Model Year = 2003-2004 (base: 2000-2002)	-1.059	0.779
Model Year = 2005-2006	-0.884	0.862
Model Year = 2007	-1.046	0.954
Model Year = 2008-2009	-0.838	0.998
Model Year = 2010	-1.330	1.259
Model Year = 2011	0.613	1.536
Honda Accord	0.231	0.163
Toyota Camry	0.098	0.177
Zip Code = 22911	-0.692***	0.230
Zip Code = 64055	-2.647**	1.185
Zip Code = 66202	-0.778***	0.153
April-May 2011	-0.203	0.198
June-Aug. 2011	0.080	0.183
Sept- Oct 2011	0.265	0.229
Nov-Dec 2011	-0.003	0.281
Std(η)	1.785**	0.894
Residual from the pricing equation	0.479	0.857
Alternative choices		
Constant	-3.046***	0.316
Log(Current Inventory-1)	1.136***	0.127
N	1585	

Note: 1) Multinomial Logit estimation of whether a car or any of its competitors of the same model is sold on a certain day. The estimation uses three Japanese models (Honda Accord, Toyota Camry and Nissan Altima) listed by CarMax in four zip codes (21162, 22911, 64055, and 66202) from January to December 2011. The second and third day of a car's listing period are used. 2) Base of Wheelbase is 103-105; base of Model year is 2000-2002; base of Time is Jan.-March 2011. 3) * $p < 0.05$, ** $p < 0.01$, *** $p < 0.1$

Table 12: Estimated Structural Parameters of the Seller's Model

Parameters	Coefficients	Std
The precision of the initial examination		
σ_0	0.301***	0.0006
Menu cost		
ϕ_1	\$315.19***	32.881
The inventory effect		
ϕ_2	-0.0287***	0.0003
N	1585	

Table 13: Policy Experiment 1: the Value of Information

(Initial inspection, Learning)	Yes	No
Yes	\$352	\$347
No	\$109	\$0

*Note:*The expected profits are computed as relative to that in the case with neither initial inspection nor learning.

Table 14: Policy Experiment 2: the Value of Lowering the Menu Cost

Menu cost	$\hat{\phi}_1/2$	$\hat{\phi}_1$	$2\hat{\phi}_1$
Expected profit	\$153	\$50	\$0

*Note:*Estimated menu cost $\hat{\phi}_1 = \$315$. The expected profits are computed as relative to that in the case of menu cost being $2\hat{\phi}_1$.

References

- ACKERBERG, D. (2003): “Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination,” *International Economic Review*, 44(3), 1007–1040.
- AGHION, P., P. BOLTON, C. HARRIS, AND B. JULIEN (1991): “Optimal Learning by Experimentation,” *Review of Economic Studies*, 58, 621–654.
- ARAMAN, V., AND R. CALDENTEY (2009): “Dynamic Pricing for Nonperishable Products with Demand Learning,” *Operations research*, 57(5), 1169.
- AVIV, Y., AND A. PAZGAL (2005): “Dynamic Pricing of Short Life-cycle Products through Active Learning,” Discussion paper, Washington Univ. in St. Louis.
- BENÍTEZ-SILVA, H., G. HALL, G. HITSCH, G. PAULETTO, S. BROOK, AND J. RUST (2000): “A Comparison of Discrete and Parametric Approximation Methods for Continuous-State Dynamic Programming Problems,” Discussion paper, Yale University.
- CHING, A., T. ERDEM, AND M. KEANE (2011): “Learning Models: An Assessment of Progress, Challenges and New Developments,” Discussion paper, University of Toronto.
- CRAWFORD, G., AND M. SHUM (2005): “Uncertainty and Learning in Pharmaceutical Demand,” *Econometrica*, 73(4), 1137–1173.
- EASLEY, AND KEIFER (1988): “Controlling a Stochastic Process with Unknown Parameter,” *Econometrica*, 56, 1945–1064.
- ERDEM, T., AND M. KEANE (1996): “Decision-making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets,” *Marketing Science*, 15, 1–20.
- ERDEM, T., M. KEANE, AND B. SUN (2008): “A Dynamic Model of Brand Choice When Price and Advertising Signal Product Quality,” *Marketing Science*, 27(6), 1111–1125.
- FERNANDEZ-VILLAVERDE, J., AND J. RUBIO-RAMIREZ (2007): “Estimating Macroeconomic Models: A Likelihood Approach,” *Review of Economic Studies*, 74, 1059–1087.
- FLURY, T., AND N. SHEPHARD (2008): “Bayesian Inference Based Only on Simulated Likelihood: Particle Filter Analysis of Dynamic Economic Models,” Discussion paper, Oxford University.
- GALLANT, R., H. HONG, AND A. KHWAJA (2009): “Estimating a Dynamic Oligopolistic Game with Serially Correlated Unobserved Production Costs,” Discussion paper, Duke University.

- GARDETE, P. (2012): “The Value of Market Information in the Dynamics of a Capital-Intensive Industry: The Case of DRAM Manufacturing,” Discussion paper, Stanford University.
- GROSSMAN, S. J., R. E. KIHLMSTROM, AND L. J. MIRMAN (1977): “A Bayesian Approach to the Production of Information and Learning by Doing,” *Review of Economic Studies*, 44, 533–547.
- HUANG, G. (2009): “Posted Price and Hagglng in the Used Car Market,” Discussion paper, John Hopkins University.
- MASON, R., AND J. VÄLIMÄKI (2011): “Learning about the Arrival of Sales,” *Journal of Economic Theory*.
- McFADDEN, D., ET AL. (1978): *Modelling the choice of residential location*. Institute of Transportation Studies, University of California.
- MIRMAN, L., L. SAMUELSON, AND A. URBANO (1993): “Monopoly Experimentation,” *International Economic Review*, 34, 549–563.
- NAIR, H. (2007): “Intertemporal Price Discrimination with Forward-Looking Consumers: Application to the US Market for Console Video-Games,” *Quantitative Marketing and Economics*, 5(3), 239–292.
- PETRIN, A., AND K. TRAIN (2010): “A Control Function Approach to Endogeneity in Consumer Choice Models,” *Journal of Marketing Research*, 47(1), 3–13.
- ROTHSCHILD, M. (1974): “A Two-armed Bandit Theory of Market Pricing,” *Journal of Economic Theory*, 9, 185–202.
- RUBIN, D. (1988): “Using the SIR Algorithm to Simulate Posterior Distributions,” in *Bayesian Statistics 3*, ed. by J. Bernardo, M. DeGroot, D. Lindley, and A. Smith. Oxford University Press.
- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999–1033.
- SUDHIR, K. (2001): “Competitive Pricing Behavior in the Auto Market: A Structural Analysis,” *Marketing Science*, pp. 42–60.
- TREFLER, D. (1993): “The Ignorant Monopolist: Optimal Learning with Endogenous Information,” *International Economic Review*, 34, 565–581.
- VILLAS-BOAS, J. (2007): “A Note on Limited versus Full Information Estimation in Non-linear Models,” Discussion paper, University of California, Berkeley.
- XU, X., AND W. HOPP (2005): “Dynamic Pricing and Inventory Control: The Value of Demand Learning,” Discussion paper, Northwestern University.