HARVARD | BUSINESS | SCHOOL

Invest in Information or Wing It? A Model of Dynamic Pricing with Seller Learning

Guofang Huang
Hong Luo
Jing Xia

Working Paper 16-027

Invest in Information or Wing It? A Model of Dynamic Pricing with Seller Learning

Guofang Huang
Carnegie Mellon University

Hong Luo
Harvard Business School

Jing Xia
Harvard University

Working Paper 16-027

Invest in Information or Wing It? A Model of Dynamic Pricing with Seller Learning*

Guofang Huang†

Hong Luo‡

Jing Xia §

First version: September 2012, Current version: September, 2015

Abstract

Pricing idiosyncratic products is often challenging because the seller, ex ante, lacks information about the demand for individual items. This paper develops a model of dynamic pricing for idiosyncratic products that features the optimal stopping structure and a seller that learns about item-specific demand through the selling 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 demand information that the dealer obtains through the initial assessment and subsequent learning in the selling process. With the dealer's average net profit per car in the estimation sample being around $1150, the initial assessment is worth around $101, and the subsequent learning in the selling process helps improve the dealer's profit by at least $269. These estimates suggest a potentially high return to taking the "information-based" approach to pricing idiosyncratic products.

Keywords: dynamic pricing, idiosyncratic products, item-specific demand, demand uncertainty, active seller learning, the value of information.

*We thank Dirk Bergemann, Guido Imbens, Przemek Jeziorski, Zhenyu Lai, Brad Larsen, Ariel Pakes, Greg Lewis, Jiwoong Shin, K. Sudhir, Juuso Valimaki and Zhixiang Zhang for helpful discussions. We also benefited from comments by seminar participants at Yale, Duke, CMU, Univ. of Wisconsin-Madison, UBC, Washington University at St Louis, USF and the Econometric Society Winter Meeting 2015. Any remaining errors are our own.

†Carnegie Mellon University, gfhuang@andrew.cmu.edu.

‡Harvard Business School, hluo@hbs.edu

§Harvard University, jingxia@fas.harvard.edu

1 Introduction

Pricing is a challenging task for dealers selling idiosyncratic products, such as used cars, houses, artwork, etc. These products show significant *item-specific* heterogeneity even after accounting for all their standard observable attributes. Furthermore, the demand for an item can also depend on the *current* preference of the *local* population. Take the example of used cars. Identical new cars can end up as used cars in very different condition after logging the same mileage, depending on how they were driven and maintained. Figure 1 shows that the Kelley Blue Book (KBB) "private party" price for a 2007 Honda Accord LX sedan with 68,500 miles, in Rockville, Maryland, ranges from \$10,400 for the "Fair condition" to \$12,550 for the "Excellent condition." Furthermore, for a particular used car in a specific condition, the KBB price can vary by market because of the differences in the local consumers' preferences.[1]

The main challenge for dealers in pricing idiosyncratic products is that they, *a priori*, often lack information on the heterogeneity of individual items or local consumers' preference for them because, by definition, they are not the prior owners of those items and often need to acquire and sell them in large numbers. To deal with this challenge, dealers may acquire more information on the demand for each item through multiple channels. They may inspect and research items they have acquired before selling them. Furthermore, they can also learn new information in the subsequent selling process by, for example, observing instances when no sale is made and communicating directly with buyers. CarMax, the largest used-car dealership in the U.S., takes such an information-based approach and makes pricing one of its biggest competitive advantages.[2]

Though potentially beneficial, the information-based approach requires costly investments (e.g., in the information and pricing system) and nontrivial fixed and variable

[1]The Kelley Blue Book is a car valuation and research company that is well known in the automotive industry. The KBB car values are based on its surveys of dealers' actual transactions. Figure 1 is a screenshot of a webpage on kbb.com from May 2012. The corresponding prices also vary by location.

[2]CarMax thoroughly inspects every car before putting it up for sale. It maintains a proprietary database of all its past transactions and uses it to analyze local consumers' preferences. The dealership tags every car on the lot with an RFID tag that tracks its location, how long it sits on the lot, and when test drives occur. Each sales person is equipped with a hand-held device in order to communicate real-time information to the central pricing department. Furthermore, CarMax's proprietary information and pricing system enables it to set and adjust its prices according to the latest information available in each stage of the selling process. Austin Ligon, the former CEO of CarMax, described its proprietary information and pricing system as "one of our biggest competitive advantages" and noted: "We adjust prices to the marketplace literally on a daily basis." Sources: (1) "CarMax Strategy Teams" (available at http://ieee.illinois.edu/wordpress/wp-content/uploads/2013/02/CarMax-Strategy-Group-Teams.pdf); and (2) "CarMax-CEO-Interview" (Mark Haines, CNBC/ Dow Jones Business Video, October 1, 2002).

costs. Should firms make the necessary investments to take such a approach to pricing idiosyncratic products? More specifically, what is the value of the information that sellers can acquire through their initial assessment and subsequent learning during the selling process? From a theoretical perspective, how does the seller's learning in the selling process affect pricing dynamics? To answer these questions, we develop a structural model of dynamic pricing with seller learning. We estimate the model using novel panel sales data from CarMax, and we use the estimated structural model to quantify the value of the demand information that the dealer acquires in the selling process.

Our theoretical model of dynamic pricing is cast as an optimal stopping problem. Specifically, we consider the problem of a seller selling an item to sequentially arriving buyers and use ξ, a continuous random variable, to describe the item-specific demand shifter. Before selling the item, the seller receives a signal—which quantifies the result of the seller's initial assessment—about ξ from a distribution centered around ξ. In the subsequent selling process, the seller receives a new signal about ξ each time a buyer decides not to buy the item. We use a static discrete-choice model to describe each buyer's purchase decision, and adopt the Bayesian Gaussian learning framework to formalize the seller's learning process. The seller's objective is to choose prices based on her latest information to maximize the present value of her expected profit from selling the item. The seller incurs a cost whenever she changes the price. The model is essentially an optimal stopping problem since the prices that the seller sets control the probabilities of sale (stopping).

We derive a few insights on pricing dynamics from our model. First, the optimal price tends to go down over time because of the dynamic adverse selection of unsold items. The fact that current buyer does not purchase the item implies that ξ was more likely overestimated. Thus, the seller is more likely to adjust her belief and, hence, the price downward after learning new information about ξ. Besides the dynamic adverse-selection effect, the dynamics in the seller's strategic responses to the learning opportunities lead to steeper price drops over time. One such effect derives from the seller's incentive to set a higher current price to delay sale and, hence, benefit from the new information. Because the value of new information drops as the seller becomes more informed, this incentive weakens over time. As a result, the optimal price can drop even when the seller's expected value of ξ remains the same. Another effect derives from the influence of the current price on the gain from subsequent learning. The current price determines the continuation probability at each value of ξ, which, in turn, affects the distribution of the new signal and the value of the new information. Thus, the seller wants to set the current price so as to gain more from learning the new information. This incentive is stronger at the beginning than later in the process, leading to additional price drops under certain conditions.

We estimate our structural model using 2011 car-level panel data from a CarMax store. The data include detailed car attributes and daily list prices set by the dealer over each car's duration on the market.[3] We estimate the demand model using the control function approach (c.f. Petrin and Train (2010)) and the structural model of dynamic pricing using the Nested Fixed Point algorithm (c.f. Rust (1987)). Estimating the dynamic pricing model is challenging because the state variables summarizing the seller's belief about ξ are not observable to us and need to be integrated out of 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 show that the information on car-specific demand that CarMax uses in its pricing is of significant value. In particular, the initial assessment increases the dealer's expected profit per car by around \$101, and the subsequent learning in the selling process improves the dealer's profit by at least \$269. These values are significant, given the net profit per car of around \$1150 in our estimation sample. Thus, the information-based approach to pricing idiosyncratic products seems worth considering, especially for large dealers that can implement it efficiently.

This paper makes the following contributions to the literature. First, to the best of our knowledge, it is the first empirical study of a dynamic-pricing problem for idiosyncratic products in which the seller's learning plays a central role. Our results shed light on the empirical importance of the seller's learning and the information-based approach in this class of pricing problems. Second, some of the insights we derive on pricing dynamics under the optimal pricing strategy are new to the related theoretical literature. Finally, besides the used-car market, which is of significant economic importance,[4] our modeling framework may also be applied to other markets of idiosyncratic products.

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3 introduces the data and presents some model-free evidence of dealers' uncertainty about car-specific demand and learning. Section 4 sets up the model and discusses the key features of the seller's optimal pricing strategy. Section 5 describes the estimation method. Section 6 presents the empirical results. Section 7 concludes.

[3] The prices at the dealer are non-negotiable.

[4] According to Ward's Automotive Yearbook 2013, about 42 million used cars (\$380 billion in total revenues) were sold in the U.S. in 2012. In comparison, about 15 million new cars (\$300 billion in total revenue) were sold in the U.S. in the same year.

2 Related Literature

This paper is related to the theoretical literature on dynamic pricing with demand uncertainty (c.f. Rothschild (1974), Grossman et al. (1977), Easley and Keifer (1988), Aghion et al. (1991), Mirman et al. (1993), Trefler (1993) and Mason and Välimäki (2011)).[5] Most of the literature is concerned with pricing problems in which the seller needs to learn about the demand for *homogeneous* products. We focus on the pricing problem for *idiosyncratic* products, which is distinguished by its inherent optimal stopping structure. Our structural model is most closely related to that of Mason and Välimäki (2011), who study the general problem of optimal stopping when the environment changes because of learning. Our model differs from theirs mainly in the description of the seller's uncertainty about demand. In their model, the uncertainty is motivated by the seller's lack of information on the buyers' arrival rate, and the true arrival rate is assumed to be either high or low. In contrast, in our model, the uncertainty derives from the seller's lack of information about item-specific heterogeneity or preference of local consumers, which is modeled as a continuous variable. Because of this difference, our model is able to capture additional pricing dynamics (e.g., those driven by the dynamics in the value of new information)[6] and it may also better capture pricing problems in which the seller's information about the condition of and the local preference for individual items is limited—such as those facing used-car dealerships or banks selling a large number of foreclosed houses,.

Riley and Zeckhauser (1983) is another closely related theoretical paper, which studies the optimal mechanism for selling an item to sequentially arriving buyers. They provide a constructive proof showing that it is optimal for the seller to charge a fixed price to, instead of haggling with, each buyer. In proving the result, they allow the seller to have imperfect information on the distribution of buyers' reservation values and to learn about them over time. Their result provides a possible explanation for why the dealer in our empirical application uses a no-haggle pricing policy. We take the policy as given and focus on the mechanisms through which the seller's learning impacts pricing dynamics and empirically quantify the value of information about item-specific

[5] 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 homogeneous products and use different approaches than ours. We are also not going into the large literature on revenue management for the same reason.

[6] The seller's belief in the model of Mason and Välimäki (2011) is described by a Bernoulli distribution. The mean and variance of the seller's belief are, thus, inter-dependent: if the mean is x, then the variance would be $x(1-x)$. The specification limits their ability to separately describe the seller's expectation and the accuracy of the seller's belief. In contrast, the seller's belief in our model is described by a normal distribution, in which the mean and variance are two independent variables.

demand.

The empirical literature on dynamic pricing for idiosyncratic products is small. A related paper from this literature is Merlo et al. (2013). Their paper builds a structural empirical model that helps explain a number of stylized facts about the pricing dynamics of individual houses, as documented by Merlo and Ortalo-Magne (2004) using rich panel data from the UK. Our paper differs from theirs in some important aspects. They focus on the problem for individuals selling their own homes, while we study the problem for dealers who are not the prior owners of the items being sold. Because of the difference in the empirical context, some of our model's key elements and the insights that we derive are different. For example, the list prices of houses in their model are subject to negotiation, whereas the used-car prices in our model are not negotiable. More importantly, the seller in their model has complete information about demand, whereas one of our main assumptions is that the seller is uncertain about demand and learns about it over time. As a result, different mechanisms (mostly due to seller learning) drive the pricing dynamics in our model.

Our paper also is related to the growing empirical literature that investigates the effects of information on the functioning of various selling mechanisms used in the secondary durable goods market (used-car market, in particular). For example, Lewis (2011) shows that disclosing verifiable information about used cars in the online retail auctions mitigates the classic adverse-selection problem. Through a large-scale field experiment, Tadelis and Zettelmeyer (2014) find that disclosing additional inspection information about used cars increases the expected revenue from the cars at wholesale auctions. Their analysis suggests a novel channel through which information disclosure affects revenue in wholesale auctions: the additional information disclosed leads to better matching of heterogeneous buyers (i.e., used-car dealers) to simultaneously auctioned cars of different qualities and, consequently, to more intense competition among buyers at each auction. In another study, Larsen (2014) estimates an alternating-offer bargaining model with two-sided incomplete information, using rich data on bargaining offers from the used-car wholesale market. The paper shows that the efficiency loss in the bargaining due to asymmetric information in this market is small. In contrast to these papers, we focus on dynamic pricing as a selling mechanism used by dealers in the off-line used-car retail market.

The broader empirical literature on dynamic pricing focuses mainly on pricing problems for homogeneous products. One related paper is Ching (2010), which studies entry by generic drugs and post-entry price competition between generic and brand-name drugs. The paper estimates a drug demand model and calibrates the manufacturers' oligopoly pricing model, in which all the manufacturers and the current patients learn about the quality of the newly introduced generic drugs through the experience of previous patients. Ching's policy experiments show that the model is able to explain the observed price increase by brand-name drugs and the price decreases by generic drugs after generic entry. Our model shares the general feature of the interaction between the seller's dynamic pricing strategy and the seller's learning process, but has very different mechanisms through which seller learning impacts pricing dynamics.[7]

The general Bayesian learning framework used in our model has been widely adopted in the literature of empirical industrial organization and marketing that studies consumers' learning behavior and its implications for demand in various consumer-goods markets. Well-known examples from this literature include Erdem and Keane (1996), Ackerberg (2003), Crawford and Shum (2005) and Erdem et al. (2008).[8] Our focus on idiosyncratic products and supply-side pricing decisions sets our paper apart from these others. The learning in our model can also be distinguished by the fact that the new information that the seller learns in each stage is endogenous to her pricing decision.

3 Data and Model-Free Analysis

3.1 Data

The data used in this paper are scraped from *Cars.com*, one of the two largest automotive classified sites in the U.S. The data cover all used cars listed by dealers in a suburban area near Baltimore in 2011. Available information includes detailed car characteristics and daily list prices for the cars' entire duration on *Cars.com*. Because cars are typically removed immediately from the website once they are sold by the listing dealers,[9] the data also allow us to determine the dates on which cars were sold.

[7]Another well-known example of this literature is Nair (2007), who shows that consumers' forward-looking behavior significantly limits the seller's ability to price discriminate inter-temporally. Our paper differs from Nair (2007) both in its substantive focus and in its modeling framework. We abstract away from consumers' forward-looking behavior in our model because, as we will explain in detail later, there is very limited room for consumers to benefit from being forward-looking in the used-car retail market, and, thus, it does not seem an important factor affecting dealers' dynamic pricing strategy.

[8]Ching et al. (2011) provide a comprehensive survey of the empirical literature on consumer learning.

[9]According to people from the industry, the main motivation for dealers to update their listings quickly is to avoid antagonizing customers who go to the stores only after identifying on the internet the cars that they want to inspect more closely. We were told that this concern is so important that CarMax temporarily suspends a car's listing whenever it is taken out for a test drive.

We focus on CarMax in our empirical analysis for the following reasons. First, as mentioned above, CarMax is a particularly good example of systematically acquiring and utilizing information in pricing decisions. This feature makes CarMax ideal for an empirical study of the value of information and learning to dealers. Second, CarMax's "no-haggle" pricing policy means that the last list prices in our data are the actual transaction prices. Furthermore, CarMax lists its entire inventory on *Cars.com* during our data period.[10] Accurate information on cars' transaction prices and brand-level inventories are important for analyzing the dealer's pricing behavior. Finally, CarMax stores often have inventory that is many times larger than that of the largest competitors in their local markets. Its dominant market position makes it less restrictive for us to abstract away from the competition from other local dealers when modeling CarMax's dynamic pricing behavior.

It is worth noting that some cars listed on *Cars.com* were eliminated from the website possibly because they were taken to wholesale auctions instead of being sold in the retail channel. These tend to be cars that have been on the market for a long time. In the case of CarMax, this issue should cause little concern because as reported in CarMax's 2011 Annual Report, "Because of the pricing discipline afforded by the inventory management and pricing system, more than 99% of the entire used car inventory offered at retail is sold at retail." In addition, it is sufficient for us to use only the first few days' data for each car to estimate our structural model. Thus, the impact of potential mismeasurement of the sale dates should be small.

In the following, we report a few stylized pricing patterns that are common to CarMax and other dealers in our data. These patterns provide some preliminary evidence suggesting that the seller is uncertain about item-specific demand and learns about it over time. We also briefly describe some differences in the pricing and sales patterns across dealers, which our pricing model might help explain.

3.2 Model-Free Analysis

Stylized Pricing Patterns

We first focus on the pricing patterns for CarMax. First, cars usually take a few days to sell, which presents enough opportunities for the seller to adjust prices. Table 1 summarizes the distribution of the time to sell. It takes the CarMax store in our data 14 days, on average, to sell a car, and 90 percent of its cars are sold within 31 days.

[10] In its 2011 annual report, CarMax says that it "lists every retail used vehicle on both *Autotrader.com* and *Cars.com*."

Second, a significant share of cars experience substantial price adjustments, most of which are decreases in price. The first two columns of Table 2 tabulate the total number of price adjustments during the cars' entire time on the market. About 30% of cars sold at the CarMax store had their listing price adjusted at least once, and the maximum number of adjustments was seven.

The top panel of Table 3a summarizes the one-time price changes—i.e., the current price relative to that on the previous day (conditional on the change being non-zero)—separately for price increases and decreases. A majority, 91.4%, of the one-time price changes are decreases. The average magnitudes are $731.9 and $499.9, respectively, for increases and decreases. The bottom panel of Table 3a summarizes the total price changes—i.e., the difference between the cars' last listing prices and their initial prices, conditional on the total price change being non-zero. Again, a majority, 92.1%, of the changes are decreases. The average magnitudes of total adjustments are $725.4 and $631.1, respectively, for increases and decreases.

Third, there is a downward trend in the conditional price-changing likelihood by cars' time on the market. Figure 2 plots the percentage of the remaining cars with prices changed relative to their prices on the previous day by time on the market. It shows that the conditional price-changing likelihood drops over the first few days, and then largely flattens out.

The above pricing patterns are not unique to CarMax. We see similar patterns for other dealers in our data. For comparison, we focus on other top dealers in this local market. Table 1 shows that these dealers also take some time, 35 days on average, to sell their cars. Table 2 shows that a large share, 46%, of cars that these dealers sell experience price adjustments. Table 3b shows that 86.5% of the one-time price changes are decreases, and the average magnitudes of the one-time adjustments are $1308.9 and $865.1, respectively, for increases and decreases; 88.5% of total price changes are decreases, and the average magnitudes of the total adjustments are $1258.9 and $1605.6, respectively, for increases and decreases. Figure 3 shows a similar decline in the conditional price-adjustment probability over time for these dealers, though the trend is less significant.

Our Hypothesis

How can one explain the pricing patterns identified above? Given that the demand for used cars was relatively stable in 2011 and that most cars are sold within a relatively short time, the price changes in our data are most likely driven mainly by two factors. One is inventory fluctuations, and the other is the seller's learning about car-specific demand.

Inventory fluctuations may lead to price adjustments because the seller may want to raise prices if the inventory drops below its desired level and lower prices if the inventory exceeds the desired level. However, inventory fluctuations should be roughly equally likely to generate price increases and decreases. Table 4 summarizes the daily percentage change in the levels of the total inventory of CarMax, the inventory of CarMax's top six car models and that of the top model of the other five top dealers. It is clear that the distributions of inventory changes are roughly symmetric around zero, which makes them unlikely to be a factor driving the systematic price decreases over time.

On the other hand, if the seller faces uncertainty about car-specific demand and learns about it over time, the dynamic adverse selection of unsold cars would generate significantly more price decreases than price increases. The intuition is that cars for which the seller overestimated the demand would be overpriced and, thus, more likely to stay on the market. Thus, the seller would be more likely to make downward adjustments in her estimates and in the prices for cars remaining on the market. As we will explain in more detail later, there are additional mechanisms through which the learning process can generate systematic price decreases over time. Lastly, the overall decline in the likelihood of price changes is consistent with dynamic pricing with menu cost and seller learning because the impact of learning on pricing diminishes as the seller's information about the remaining cars improves over time.[11]

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 the price sequentially—that is, to skim the price-inelastic consumers first and sell to the price-elastic consumers later. However, this does not seem a compelling story for the used-car market (and the secondary durable goods market more broadly). Each used car is a unique product and normally is sold within a few days. In addition, most cars (around 70% for CarMax) are sold without ever having their prices changed. So, if a buyer is interested in buying a particular car, he very likely will miss the opportunity to buy the car if he waits for the price to drop. Therefore, the forward-looking behavior of consumers seems to have limited relevance

[11]Conversations with people with work experience in the pricing department of CarMax also confirmed the importance of information collection and learning in the price-adjustment decisions.

in this market and, thus, is unlikely to explain the pricing patterns described above.[12]

Alternatively, time on the market may work as a signal of car quality to consumers. In particular, consumers may have heterogeneous imperfect information about the quality of used cars. In this case, one reason that a car has not been sold is that the previous buyers had unfavorable information about its quality. Then, a longer time on the market can signal worse car quality (c.f. Taylor (1999)). If the signaling effect is significant, both the car-specific demand and price would decrease with time on the market. Our empirical evidence, however, does not support the signaling effect as a major driving force of the price patterns. Although consumers can filter cars listed on *Cars.com* by their time on the market (up to a few ranges of days), but a sample of consumer browsing data that we obtained from the website show that less than one percent of the searches used the option. Our conversations with industry insiders also corroborate that, empirically, time on the market is unlikely to be an important factor affecting consumers' evaluation of used cars.

Some Differences in the Pricing Patterns across Dealers

The data show some substantial differences in the pricing patterns and sales performance across dealers, especially when comparing CarMax to other dealers. For example, compared to other dealers, CarMax takes a significantly shorter time to sell its cars (c.f. Table 5a) and adjusts prices for a smaller share of its cars (c.f. Table 2). In addition, the total price adjustments made by CarMax are, on average, smaller than those made by other dealers (c.f. Table 5b). Though it is not our aim to systematically explain these across-dealer differences, our analysis in the following sections suggests that these differences may be driven partly by the variation in the dealers' ability to examine and research their cars and to incorporate new information learned in the selling process into their pricing decisions.

Some other factors that this paper does not focus on may also help explain the above differences. For example, dealers' inventory and brand portfolio sizes may be relevant. Dealers carrying a larger inventory and more brands would attract more consumers, which may allow them to sell their cars faster and to have less need to adjust prices. Consistent with this explanation, CarMax's inventory is larger and covers a more comprehensive list of brands. However, the particular set of brands that each dealer carries does not seem very relevant. Table 6 summarizes the time to sell

[12]Related to our observations here, Sweeting (2012) finds little evidence for strategic consumer purchase behavior in the secondary market for Major League Baseball tickets, and he shows that a dynamic pricing model with static time-invariant demand fits his data well.

and total price changes for the top six models at CarMax. The table shows that the pricing behavior and sales performance of these models are similar to each other, which is inconsistent with pricing and sale performance being strongly car-brand dependent.

In summary, the evidence presented in this section suggests that in order to explain the pricing patterns observed in the data, it is essential to incorporate the seller's uncertainty and learning about item-specific demand. In the following, we develop a structural model of dynamic pricing with these features and apply it to the CarMax data. The lessons from the exercise are also valuable for understanding the dynamic pricing problems of idiosyncratic products in general.

4 Model

4.1 Model Overview

The context of our model is a used-car retail market with a monopolist dealer. To fix ideas, let us consider the seller's problem of setting non-negotiable prices for a particular car when she faces sequentially arriving buyers. The seller is uncertain about the demand for the car. To reduce the uncertainty, she may first inspect the car and evaluate local consumers' preferences for the car. Based on the initial assessment, the seller sets the price for the first buyer. In the subsequent selling process, she receives additional information about the demand every time a buyer decides not to buy the car. Each buyer has a demand for, at most, one car. After a buyer makes a one-shot purchase decision, he exits the market and never returns. The seller sets a price for the car based on her latest information before the arrival of each subsequent buyer.

Two assumptions that we make in our model worth clarifying here. First, we assume that the seller considers the pricing problem for each car separately. This assumption is motivated by the fact that, for idiosyncratic products like used cars, the seller's uncertainty and learning about demand is mostly car-specific. In addition, it would be computationally intractable to explicitly model the dynamic pricing problem when jointly maximizing the profit of multiple cars, and this can also be difficult for dealers in reality.

Second, we assume that one and only one buyer arrives each day. This normalization assumption is necessary because we do not observe any information about the buyers in our data. Thus, we are essentially modeling the daily demand for each individual car, and the seller learns about the uncertain component of it.

4.2 Demand Model

We use a static discrete-choice model to describe the demand for each car. Consider the demand for car j on day t (when it is still available). Let X_j be a vector of observable attributes of car j and ξ_j a scalar that summarizes all other factors that affect the buyers' average valuation for the car. We refer to ξ_j as the car's (latent) quality, assuming that it is observable to all buyers but not observable to the seller. Let v_{jt} be buyer t's value for purchasing car j at the price of p_{jt}. We specify v_{jt} as follows:

$$v_{jt} = X_j\beta + \alpha p_{jt} + \xi_j + \varepsilon_{jt}, \tag{1}$$

where ε_{jt} is buyer t's idiosyncratic preference shock for car j, and β and α are, respectively, the marginal value for X_j and the price. Furthermore, let buyer t's value of buying cars other than car j be v_{-jt}, and let the value of not buying any car be v_{0t}. We specify the values of these two choices as follows:

$$\begin{aligned} v_{-jt} &= u_{-jt} + \varepsilon_{-jt}, \\ v_{0t} &= \varepsilon_{0t}, \end{aligned}$$

where u_{-jt} is the mean value of buying other cars; the mean value of not buying any car is normalized to zero; and ε_{-jt} and ε_{0t} are buyer t's idiosyncratic preference shocks for the corresponding choices. In our empirical application, we approximate u_{-jt} using a function of the number of other cars in car j's segment on the same day, K_{jt}:

$$u_{-jt} = \bar{u}_{-j} + \phi_0 \log(K_{jt}),$$

where $\bar{u}_{-j}$ is the choice-specific constant.[13] Thus, we assume that for a buyer considering whether to buy car j, his alternatives are restricted to cars in the same segment as car j. This simplification is motivated partly by the need to keep the demand-side model parsimonious so that the corresponding dynamic pricing model is computationally tractable.

Define I_{jt} as an indicator function of buyer t choosing to buy car j:

$$I_{jt} = 1\left\{v_{jt} > v_{-jt} \;\&\; v_{jt} > v_{0t}\right\}.$$

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

[13] The approximation may be interpreted as aggregating the choices of buying other cars as a single option (McFadden et al. (1978)).

Similarly, we define I_{-jt} as follows:

$$I_{-jt} = 1\left\{v_{-jt} > v_{jt} \ \& \ v_{-jt} > v_{0t}\right\}.$$

Finally, we assume that, for all t, the preference shocks $\varepsilon_t \equiv (\varepsilon_{jt}, \varepsilon_{-jt}, \varepsilon_{0t})$ are drawn from the same mulitvariate normal distribution and are independent of ξ_j.

4.3 Dynamic Pricing Model

In the following, we first introduce the key components of our model and then formalize it. The seller's objective is to maximize the present value of the expected profit from selling the car.

Seller Learning

The seller observes X_j but is uncertain about ξ_j, which captures her imperfect information about the demand for car j. We adopt the Bayesian Gaussian learning framework to model the seller's learning process. We assume that $\xi_j \sim N\left(0, \sigma_\xi^2\right)$. The seller can learn about ξ_j through two channels. First, she can assess the demand for the car before setting the price for the first buyer. We quantify the result of the initial assessment as an unbiased signal, y_{j0}, drawn from $N\left(\xi_j, \sigma_0^2\right)$, where the inverse of σ_0^2 captures the thoroughness of the assessment. With y_{j0}, the seller updates her belief about ξ_j using Bayes' rule.

Second, the seller can further learn about ξ_j in the selling process. This includes learning by simply observing instances when no sale is made and communicating directly with buyers. To approximate learning from these sources, we assume that the seller receives a signal $y_{jt} \equiv \xi_j + \epsilon_{jt}$ after buyer t decides not to buy car j, where $\epsilon_{jt} \sim N\left(0, \sigma_s^2\right)$ and is independent of ξ_j and ε_{jt}. The seller updates her belief about ξ_j using Bayes' rule every time she receives a new signal.[14]

For later reference, we use $y^t \equiv (y_0, ..., y_{t-1})$ to denote the vector of signals that the seller receives before the arrival of buyer t, and $\left(\mu\left(y^t\right), \sigma_t^2\right)$ to denote the mean and variance of the seller's posterior belief about ξ_j after observing y^t. To keep track of

[14] Alternatively, one may specify the signal that the seller receives after a buyer decides not to buy the car as $y_{jt} \equiv \xi_j + \varepsilon_{jt}$; that is, the seller learns the buyer's exact value for the car. Though this alternative specification completely captures learning about ξ_j from observing instances when no sale is made (since $\xi_j + \varepsilon_{jt}$ is a sufficient statistic for such an event), it lacks the empirical flexibility for the "amount" of information that the seller actually learns each day. This is because the ratio of the variance of ξ_j to that of ε_{jt} not only determines the amount of information the seller learns each day, but also determines the extent of selection (i.e., when fixing the price, cars staying on the market longer tend to be cars with lower values of ξ_j) that we observe on the demand side.

the seller's belief, we need to know only $(\mu(y^t), \sigma_t^2)$ because both the prior beliefs and the signals have normal distributions. For simplicity, we sometimes write μ_t in place of $\mu(y^t)$.

Menu Cost

The seller pays some costs—the menu cost—every time she changes the price. The most direct cost is that of updating the prices posted on cars and in the advertisements. The menu cost could also include the cost of coming up with a new price given the updated belief about ξ_j. We use a random variable, φ_{jt}, to capture the cost of changing the price for car j before buyer t arrives, and we assume that φ_{jt} is drawn from the exponential distribution with mean ϕ_1. In addition to capturing the variability in the cost of changing prices, the specification also guarantees that the seller's value function is smooth in its arguments, making it easier to numerically solve the seller's optimal pricing problem.[15] We assume that the seller observes the realized value of the menu cost φ_{jt} before she decides whether to update the price for the incoming buyer t.

Inventory Management and Competition

When selling a car, the seller incurs an "opportunity cost" determined by her current inventory level. Recall that, in the demand model, we defined K_{jt} as the number of other cars in car j's segment on day t. We specify the opportunity cost as $m(K_{jt};\phi_2)$ and assume that K_{jt} evolves as an exogenous first-order Markov process. We use this specification as a parsimonious way to capture the impact of the seller's inventory management on her pricing strategies for individual cars. Intuitively, when the current inventory is high, the chances of future stock-out is low. Then, the seller may want to set prices lower to sell faster. When the current inventory is relatively low, however, the seller may want to set prices higher to balance the current profit and potential future loss due to stock-out.

We abstract away from competition from the seller's competitors because we focus on a dominant dealer in our empirical application. Our specification, however, does capture the dealer's own inventory competing for the same current demand, and the price that the seller sets for a car would respond to such a competition effect.

[15] Smooth value functions can be accurately approximated, as needed in the numerical solution of the dynamic pricing model. The value function would have kinks if the menu cost is assumed to be fixed, and functions with kinks are difficult to approximate accurately.

Holding Cost

The seller also pays a cost for holding on to the car for another day. The holding cost, assumed to be a constant, ϕ_3, includes the cost of maintaining the car (e.g., having salespeople help with test drives, cleaning up cars after test drives, keeping the car filled with gasoline, etc.) until it is sold. It can also include the opportunity cost of the occupied parking space when the seller operates at capacity.

In the following, we formalize the dynamic pricing model featuring the above elements. We use the function $D(p_{jt}, K_{jt}, \xi_j) \equiv E_{\varepsilon_t} I_{jt}(X_j, p_{jt}, K_{jt}, \xi_j, \varepsilon_t)$ to denote buyer t's probability of buying car j conditional on (p_{jt}, K_{jt}, ξ_j), emphasizing its dependence on price p_{jt}, inventory K_{jt} and the car's quality ξ_j, while suppressing the dependence on the time-invariant variables X_j for simpler notation. In what follows, we also suppress the car index j.

Let $p_t : R^{t+3} \to R^+$ be a pricing function that maps the vector of state variables, $(y^t, K_t, p_{t-1}, \varphi_t)$, to a price. Then, the seller's pricing strategy can be described as $(p_t)_{t=1}^{\infty}$, which maps the seller's latest information at the beginning of every day to a price.[16] Formally, the seller's profit-maximization problem can be written as follows:

$$\max_{(p_t)_{t=1}^{\infty}} E_{\xi} E_{(y^t)_{t=1}^{\infty}|\xi} E_{(K_t)_{t=2}^{\infty}|K_1} \sum_{t=1}^{\infty} \delta^{t-1} \chi_t E_{\varphi_t} \pi_t (p_t, \xi, K_t, \varphi_t),$$

$$\begin{aligned}
\text{s.t. } \pi_t &= (p_t - m(K_t)) D(p_t, K_t, \xi) - \phi_3, \text{ if } t = 1, \\
\pi_t &= -\varphi_t \cdot 1\{p_t \neq p_{t-1}\} + (p_t - m(K_t)) D(p_t, K_t, \xi) - \phi_3, \text{ if } t \geq 2, \\
m(K_t) &= \left(\frac{2 \exp\left(\phi_2 \left(K_t - \bar{K}\right)\right)}{1 + \exp\left(\phi_2 \left(K_t - \bar{K}\right)\right)} - 1 \right) \bar{c}, \\
\chi_t &= \prod_{\tau=1}^{t-1} (1 - I_\tau), \\
y^{t+1} &= \left(y^t, y_t\right), \\
y_t &= \xi + \epsilon_t,
\end{aligned}$$

where χ_t indicates the availability of the car at the beginning of day t; δ is the seller's discount factor; $m(K_t)$ is the opportunity cost to the seller when selling the car, which equals zero under the null hypothesis of $\phi_2 = 0$; and ϕ_3 is the daily holding cost. The opportunity cost $m(K_t)$ is determined by the deviation of the seller's current inventory from its mean level $(\bar{K})$, and it belongs to the interval of $[-\bar{c}, \bar{c}]$, where $\bar{c}$ is a constant.

[16] With a bit abuse of notation, we also use p_t to denote the value of the pricing function at a particular vector of state variables. It should be clear what the notation means in a given context.

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 (see Appendix A for details). Let us define the following value function:

$$V^t(S_t) = \max_{(p_\tau)_{\tau=t}^{\infty}} E_{\xi|y^t} E_{(y^\tau)_{\tau=t+1}^{\infty}|\xi,y^t} E_{(K_\tau)_{\tau=t+1}^{\infty}|K_t} \sum_{\tau=t}^{\infty} \delta^{\tau-t} \chi_\tau E_{\varphi_\tau} \pi_\tau(p_\tau, \xi, K_\tau, \varphi_\tau),$$

where $S_t \equiv (y^t, p_{t-1}, K_t)$. Then, the seller's profit-optimization problem has the following Bellman Equation representation:

$$\begin{aligned} V^t(S_t) = E_{\varphi_t} \max_{p_t} \{ & E_{\xi|y^t} \pi_t(p_t, \xi, K_t, \varphi_t) + \\ & E_{\xi|y^t}(1 - D(p_t, K_t, \xi)) \delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} V^{t+1}(S_{t+1}) \}, \end{aligned}$$

$$\begin{aligned} s.t.\ \pi_t &= -\varphi_t \cdot 1\{p_t \neq p_{t-1}\} + (p_t - m(K_t)) D(p_t, K_t, \xi) - \phi_3, \\ m(K_t) &= \left(\frac{2\exp(\phi_2(K_t - \bar{K}))}{1 + \exp(\phi_2(K_t - \bar{K}))} - 1 \right) \bar{c}, \\ y^{t+1} &= (y^t, y_t), \\ y_t &= \xi + \epsilon_t. \end{aligned}$$

Let us assume that the optimal pricing strategy is stationary so that it depends on y^t only through $(\mu(y^t), \sigma_t^2)$. Then, we have the following slightly more concise Bellman equation representation for the seller's profit maximization problem:

$$\begin{aligned} V(S_t) = E_{\varphi_t} \max_{p_t} \{ & E_{\xi|(\mu(y^t),\sigma_t^2)} \pi_t(p_t, \xi, K_t, \varphi_t) + \\ & E_{\xi|(\mu(y^t),\sigma_t^2)}(1 - D(p_t, K_t, \xi)) \delta E_{\mu(y^{t+1})|\xi,(\mu(y^t),\sigma_t^2)} E_{K_{t+1}|K_t} V(S_{t+1}) \}, \end{aligned}$$

$$\begin{aligned} s.t.\ \pi_t &= -\varphi_t \cdot 1\{p_t \neq p_{t-1}\} + (p_t - m(K_t)) D(p_t, K_t, \xi) - \phi_3, \\ m(K_t) &= \left(\frac{2\exp(\phi_2(K_t - \bar{K}))}{1 + \exp(\phi_2(K_t - \bar{K}))} - 1 \right) \bar{c}, \\ \mu_{t+1} &= \frac{\sigma_t^2 y_t + \sigma_s^2 \mu_t}{\sigma_t^2 + \sigma_s^2}, \\ \sigma_{t+1}^2 &= \frac{\sigma_t^2 \sigma_s^2}{\sigma_t^2 + \sigma_s^2}, \\ y_t &= \xi + \epsilon_t, \end{aligned}$$

where $S_t \equiv ((\mu(y^t), \sigma_t), p_{t-1}, K_t)$, and the value function depends on y^t only through $(\mu(y^t), \sigma_t^2)$, 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 problem with learning. By setting prices, the seller controls the probability of stopping (i.e., sale) given her latest information about ξ.

4.4 The Optimal Pricing Strategy

The above Bellman-equation representation clarifies the seller's trade-offs in her pricing decisions. The chosen price p_t determines not only the expected payoff on the current day, but also the probability $1 - D(p_t, K_t, \xi)$, with which the car stays unsold and the seller receives the corresponding option value of selling to future buyers, $\delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} V(S_{t+1})$.[17] Therefore, the trade-off between the expected current payoff and the continuation value (i.e., the option value weighted by the continuation probability) determines the optimal price. Next, we discuss the pricing dynamics—driven mainly by the seller's learning in the selling process—under the optimal pricing strategy.

Pricing Dynamics in the Presence of Seller Learning

The seller's learning impacts the price dynamics under the optimal pricing strategy via multiple mechanisms. First, learning changes the seller's belief about ξ, which directly affects the expected current demand and option value and, thus, the optimal price. Note that, in our case, learning happens only when a car is not sold. Because cars for which the seller overestimated demand (i.e., $\mu_t > \xi$) are more likely to remain on the market, subsequent learning tends to result in more pessimistic beliefs about ξ.[18] Therefore, learning coupled with the dynamic adverse selection of unsold cars by the estimate error (i.e., $\mu_t - \xi$) creates a downward trend in the optimal prices.

Second, the dynamics in the value of new information generate additional intertemporal drops in the optimal prices. The option value of selling to future buyers, $\delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} V(S_{t+1})$, derives partly from new information, y_t. Thus, the value

[17]Note that both the continuation probability and expected option value depend on ξ, and that, given p_t, the option values with higher ξ are received with smaller probabilities.

[18]This claim is straightforward to prove. First, note that at the beginning of the first day, we have that $Pr(\mu_1 > \xi) = \frac{1}{2}$. Let $R_t \equiv 1 - I_t$, denote the event that the car is not sold on day t. It follows that $Pr(R_1|\mu_1 > \xi) > Pr(R_1|\mu_1 < \xi)$ because, for any given ξ, the optimal price is higher when $\mu_1 > \xi$ than when $\mu_1 < \xi$. Then, using Bayes' Theorem, it is straightforward to verify that $Pr(\mu_1 > \xi|R_1) > \frac{1}{2}$, meaning that, more often than not, the seller's initial assessment over-estimated the demand for a car if the car is not sold on the first day. By induction, we can verify that $Pr(\mu_t > \xi|R_t) > \frac{1}{2}$, for all $t > 1$. Note that $\mu_{t+1} - \mu_t = \frac{\sigma_t^2}{\sigma_t^2 + \sigma_s^2}(y_t - \mu_t)$. Hence, we have $Pr(\mu_{t+1} - \mu_t < 0|R_t) > \frac{1}{2}$ and $E(\mu_{t+1} - \mu_t|R_t) < 0$ because $y_t \sim N(\xi, \sigma_s)$ and μ_t is also a random variable with the normal distribution.

of new information leads to higher option values and higher optimal current prices (relative to the scenario that has no new information available on day t but is otherwise the same as in our model).[19] However, because the value of new information diminishes as the seller becomes more informed, the optimal price goes down over time, ceteris paribus.

Finally, the influence of the current price on the gain from subsequent learning leads to additional dynamics in the optimal prices, which we call the "active-learning" effect. First, because the option value, $\delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} V(S_{t+1})$, increases with the new signal y_t and the mean of y_t is ξ, the seller has an incentive to set the price so as to increase the continuation probability at larger values of ξ relative to that at smaller values of ξ.[20] Second, the value of new information, as a component of the option value, depends on ξ. Thus, the seller also has an incentive to set the price to increase the continuation probability at values of ξ that are associated with greater values of new information. In particular, if the value of new information increases with ξ, the seller would want to also increase the continuation probability at larger values of ξ. It is easy to verify that these incentives to increase the continuation probability at larger values of ξ lead to a higher optimal current price if $\frac{\partial^2 D(p_t,K_t,\xi)}{\partial p_t \partial \xi} < 0$ for all ξ. Additional dynamics in the optimal price arise as these incentives weaken over time. In general, the exact price dynamics that the active-learning effect creates depend on the shape of the demand function $D(p_t, K_t, \xi)$ and how the value of new information varies with ξ.

If the seller's learning affects prices only through the first mechanism discussed above, the magnitude of the total price change over time can be a good proxy for the extent of the subsequent learning in the selling process (which, in turn, informs us about the quality of the seller's information about ξ right after the initial assessment). However, as explained above, learning also affects price dynamics through the two additional mechanisms because of the seller's strategic pricing responses to the learning opportunities. These additional mechanisms highlight the importance of developing a full-fledged model of dynamic pricing in order to quantify the value of the seller's initial assessment of the car-specific demand.

[19] Suppressing the adverse-selection effect, we can express the value of new information as $\left(E_{y^{t+1}|y^t} V\left(\mu\left(y^{t+1}\right), \sigma_{t+1}, K, p\right) - V\left(\mu\left(y^t\right), \sigma_t, K, p\right)\right)$, which is always positive and decreases over time. The positive measure of the value of new information is intuitive given the well-known Blackwell's Theorem (c.f. Blackwell (1953)).

[20] Mason and Välimäki (2011) make a similar point and call it the "controlled-learning" effect. They show that such an effect results in a higher optimal current price than in the case without learning.

Overall Patterns of Price Dynamics

We make three observations about the overall patterns of price dynamics under the optimal dynamic pricing strategy. First, the optimal price for a car tends to drop over time. The dynamic adverse-selection effect and the diminishing value of new information both create downward trends in the optimal prices. The active-learning effect may generate additional downward price shifts under certain conditions. Second, an individual price sequence can go either up or down as a result of learning and changes in inventory. Even though the seller's belief about ξ is more likely to adjust downward because of the dynamic adverse-selection effect, it is possible for the seller to adjust her belief upward when the signal about ξ is sufficiently high. This can happen, for example, when a buyer decides not to buy the car only because his valuation of the alternatives is even higher. Third, prices tend to change more frequently earlier in the selling process than later because the impact of learning is larger initially (and, hence, more likely to justify the menu cost).

5 Estimation and Identification

In this section, we describe in detail the estimation and identification of our structural models.

5.1 Estimating the Demand Model

There are two main challenges in the estimation of the demand model: (1) the price is potentially endogenous because the seller has some information about ξ_j when she sets each price; and (2) we need to estimate the distribution of the unobserved heterogeneity ξ_j. We adapt the control function approach (c.f. Petrin and Train (2010)) to deal with the issue of price endogeneity and use the observed inter-temporal pattern in the sale probabilities to identify the variance of ξ_j.

Let $Z_j \equiv \left(X_j, \widetilde{X}_j\right)$ be the exogenous variables, among which $\tilde{X}_j$ is excluded from buyers' utility function for car j.[21] Suppose that we have the following reduced-form pricing equation for the initial price:

$$p_{j1} = Z_j \varphi + \zeta_{j1}.$$

[21]The constant term is the first element in X_j, and its coefficient is β_0.

Assume that $(\xi_j, \zeta_{j1}) \perp X_j$, where $(\xi_j, \zeta_{j1}) \sim N\left(0, \begin{pmatrix} \sigma_\xi^2 & \sigma_{\xi\zeta}^2 \\ \sigma_{\xi\zeta}^2 & \sigma_\zeta^2 \end{pmatrix}\right)$. Note that $\sigma_{\xi\zeta}^2 \neq 0$ is the cause of price endogeneity. Then, it follows that $\xi_j = \frac{\sigma_{\xi\zeta}^2}{\sigma_\zeta^2}\zeta_{j1} + \eta_j$, where $\eta_j \perp (\zeta_{j1}, \varepsilon_t, p_{j1}, X_j)$, $\eta_j \sim N\left(0, \sigma_\eta^2\right)$ and $\sigma_\eta^2 \equiv \sigma_\xi^2 - \frac{\sigma_{\xi\zeta}^4}{\sigma_\zeta^2}$. This expression of ξ_j allows us to rewrite buyer 1's value for buying car j as follows:

$$v_{j1} = X_j\beta + \alpha p_{j1} + \psi\zeta_{j1} + \eta_j + \varepsilon_{j1}, \tag{2}$$

where $\psi = \frac{\sigma_{\xi\zeta}^2}{\sigma_\zeta^2}$.

Without loss of generality, let $\varepsilon_{0t} = 0$ because all that matters for the likelihood of sales data is the joint distribution of $\varepsilon_{jt} - \varepsilon_{0t}$ and $\varepsilon_{-jt} - \varepsilon_{0t}$. Denote the variance of ε_{jt} as σ_j^2, the variance of ε_{-jt} as σ_{-j}^2 and the covariance of ε_{jt} and ε_{-jt} as σ_{-jj}^2 (note that σ_j^2, σ_{-j}^2 and σ_{-jj}^2 are assumed to be independent of j). For identification, we normalize σ_j to be one. The parameters we need to estimate for the demand model are, therefore, $\theta_d \equiv (\beta, \alpha, \bar{u}_{-j}, \phi_0, \psi, \sigma_{-j}, \sigma_{-jj}, \sigma_\xi)$.

We use a two-step method to estimate the demand model. In step one, we use the first day's data for each car to estimate $\theta_{d1} \equiv (\beta, \alpha, \bar{u}_{-j}, \phi_0, \psi, \sigma_{-j}, \sigma_{-jj})$—i.e., all the demand-model parameters except for σ_ξ. We first estimate the reduced-form pricing equation to obtain the estimate of residual ζ_{j1}. We use the first day's inventory level K_{j1} as the excluded variable, following Berry et al. (1995). The inventory level is a valid excluded variable because it affects p_{j1} due to the competition effect, but does not directly affect buyer 1's utility for car j. Then, we estimate the demand as a multinomial probit model with three choices. The sale probability of car j on the first day, h_{j1}, can be written as follows:

$$\begin{aligned} h_{j1} &\equiv \Pr\left(I_{j1} = 1 | X_j, p_{j1}, K_{j1}, \hat{\zeta}_{j1}; \theta_{d1}, \sigma_\eta^2\right), \\ &= \int 1\{v_{j1} > v_{-j1}\} \cdot 1\{v_{j1} > 0\}\, dP\left(\eta_j + \varepsilon_{j1}, \varepsilon_{-j1}\right), \end{aligned}$$

where $P\left(\eta_j + \varepsilon_{j1}, \varepsilon_{-j1}\right)$ is the probability measure of the bivariate normal distribution with the mean as zeros and the covariance matrix as $\begin{pmatrix} 1 + \sigma_\eta^2 & \sigma_{-jj}^2 \\ \sigma_{-jj}^2 & \sigma_{-j}^2 \end{pmatrix}$ (note that $\eta_j \perp (\varepsilon_{j1}, \varepsilon_{-j1})$ by assumption).[22] We define h_{-j1} similarly as the probability of some other car in car j's segment being sold on car j's first day. Note that the above expressions of h_{j1} and h_{-j1} use the expression of v_{j1} in equation (2). Following Kamakura

[22]It is worth pointing out that, as the car price changes for some cars over time, the joint distribution of ξ_j and ζ_{jt} also changes. So, we focus on the first day for our estimation in the first step.

(1989), we use Mendell-Elston's analytical approximation for the computation of the above choice probabilities. Denote $\tilde{\theta}_{d1} \equiv \frac{\theta_{d1}}{\sqrt{1+\sigma_\eta^2}}$, and we use the Maximum Likelihood Estimation (MLE) method to obtain the estimate of $\tilde{\theta}_{d1}$ as follows:

$$\hat{\tilde{\theta}}_{d1} = \arg\max_{\tilde{\theta}_{d1}} \sum_{j=1}^{J} \log \left(h_{j1}^{I_{j1}} h_{-j1}^{I_{-j1}} \left(1 - h_{j1} - h_{-j1}\right)^{1-I_{j1}-I_{-j1}} \right),$$

where J is the total number of cars in our estimation sample. For later use, we denote $\hat{\theta}_{d1} \equiv \hat{\tilde{\theta}}_{d1}\sqrt{1+\sigma_\eta^2}$.

In the second step, we use the first T days' data of each car to estimate σ_ξ. Let $\tau_j \in \{1, 2, 3, ...\}$ be the number of days that car j took to sell, and define $T_j \equiv \min\{T, \tau_j\}$. Let $\tilde{h}_{jt}$ be the conditional sale probability of car j given η_j (and the observable variables) on day $t \leq T_j$. Then:

$$\begin{aligned} \tilde{h}_{jt} &\equiv \Pr\left(I_{jt} = 1 | X_j, p_{jt}, K_{jt}, \hat{\zeta}_{jt}, \eta_j; \hat{\theta}_{d1} \right), \\ &= \int 1\{v_{jt} > v_{-jt}\} \cdot 1\{v_{jt} > 0\}\, dP\left(\varepsilon_{jt}, \varepsilon_{-jt}\right), \end{aligned}$$

where $P(\varepsilon_{jt}, \varepsilon_{-jt})$ is the probability measure of the bivariate normal distribution with the mean as zero and the covariance matrix as $\begin{pmatrix} 1 & \sigma_{-jj}^2 \\ \sigma_{-jj}^2 & \sigma_{-j}^2 \end{pmatrix}$. Define $\tilde{h}_{-jt}$ similarly as the probability of some other car in car j's segment being sold on car j's day $t \leq T_j$. Then, we have the following expression for the likelihood of the observation of car j:

$$\begin{aligned} & L\left(\left(I_{jt}, I_{-jt}\right)_{t=1}^{T_j} | X_j, p_{jt}, K_{jt}, \hat{\zeta}_j; \hat{\theta}_{d1}, \sigma_\eta^2 \right) \\ = & \int \Pi_{t=1}^{T_j} \tilde{h}_{jt}^{I_{jt}} \tilde{h}_{-jt}^{I_{-jt}} \left(1 - \tilde{h}_{jt} - \tilde{h}_{-jt}\right)^{1-I_{jt}-I_{-jt}} d\Phi\left(\eta_j / \sigma_\eta\right), \end{aligned}$$

where Φ is the probability measure of the standard normal distribution, and we integrate out η_j to get the observable likelihood. We then estimate σ_η by using MLE:

$$\hat{\sigma}_\eta = \arg\max_{\sigma_\eta} \sum_{j=1}^{J} \log L\left(\left(I_{jt}, I_{-jt}\right)_{t=1}^{T_j} | X_j, p_{jt}, K_{jt}, \hat{\zeta}_j; \hat{\theta}_{d1}, \sigma_\eta^2 \right).$$

The identification of σ_η is straightforward. Larger σ_η implies more selection on η and, ceteris paribus, a faster decrease in the sale probability by day. Thus, the rate at which the conditional sale probability drops over time provides the information for identifying σ_η. We obtain the estimate of σ_ξ as $\hat{\sigma}_\xi = \sqrt{\hat{\psi}^2 \hat{\sigma}_\zeta^2 + \hat{\sigma}_\eta^2}$. The demand model parameterized by $\hat{\theta}_d = \left(\hat{\hat{\theta}}_{d1}\sqrt{1+\hat{\sigma}_\eta^2}, \hat{\sigma}_\xi^2\right)$ is used in the estimation of our dynamic pricing model as described in the following.

5.2 Estimating the Dynamic Pricing Model

In general, we use the Maximum Partial Likelihood Estimation (MPLE) method to estimate the structural parameters, $\theta_s \equiv (\sigma_0, \sigma_s, \phi_1, \phi_2, \phi_3)$, in the model.[23] We use the first T days' data of each car, as in the second step of the demand estimation, to estimate the pricing model. Limiting the number of days used in the estimation does not affect identification, but helps reduce computational burden.[24]

First note that we can write the likelihood of the observed price sequence and the sale status of car j as follows:

$$\begin{aligned} l\left((p_{jt}, I_{jt})_{t=1}^{T_j}\right) &= l\left(p_{j1}, I_{j1}\right)\prod_{t=2}^{T_j} l\left((p_{jt}, I_{jt}) \,|\, (p_{j\tau}, I_{j\tau})_{\tau=1}^{t-1}\right), \\ &= l\left(p_{j1}\right)\cdot l\left(I_{j1}|p_{j1}\right)\prod_{t=2}^{T_j} l\left(p_{jt}|\, (p_{j\tau}, I_{j\tau})_{\tau=1}^{t-1}\right) l\left(I_{jt}|\, (p_{j\tau}, I_{j\tau})_{\tau=1}^{t-1}, p_{jt}\right), \end{aligned}$$

where we suppress the dependence of the likelihood on model parameters (and the exogenous covariates). Dividing the above likelihood by $l\left(I_{j1}|p_{j1}\right)\prod_{t=2}^{T_j} l\left(I_{jt}|\, (p_{j\tau}, I_{j\tau})_{\tau=1}^{t-1}, p_{jt}\right)$, which does not depend θ_s, yields the partial likelihood of the observed price sequence of car j, $\tilde{l}\left((p_{jt}, I_{jt})_{t=1}^{T_j}|\theta_s\right)$:

$$\begin{aligned} \tilde{l}\left((p_{jt}, I_{jt})_{t=1}^{T_j}|\theta_s\right) &= l\left(p_{j1}|\theta_s\right)\prod_{t=2}^{T_j} l\left(p_{jt}|\, (p_{j\tau}, I_{j\tau})_{\tau=1}^{t-1}; \theta_s\right), \\ &= l\left(p_{j1}|\theta_s\right)\prod_{t=2}^{T_j}\int l\left(p_{jt}|y_j^t, p_{j,t-1}; \theta_s\right) f\left(y_j^t|\, (p_{j\tau}, I_{j\tau})_{\tau=1}^{t-1}; \theta_s\right) dy_j^t. \end{aligned}$$

The expression after the second equality above makes it explicit that a) the seller's optimal pricing strategy depends on her latest information, y_j^t, as well as on the price on the previous day $p_{j,t-1}$; and b) we, as researchers, do not observe the information the seller receives and, thus, y_j^t has to be integrated out. Thus, assuming that we know how to compute the integration in the above likelihood function, we now can estimate

[23] As explained later, we calibrate the discount factor δ to match an annual rate of 25%.

[24] In 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 in the estimation sample. The computation is costly, especially because the model-predicted optimal price for each day is different for different cars.

θ_s using the MPLE method as follows:

$$\hat{\theta}_s = \arg\max_{\theta_s} \sum_{j=1}^{J} \left[\log l\left(p_{j1}|\theta_s\right) + \sum_{t=2}^{T_j} \log \int l\left(p_{jt}|y_j^t, p_{j,t-1}; \theta_s\right) f\left(y_j^t | \left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-1}; \theta_s\right) dy_j^t \right]. \tag{3}$$

Following Rust (1987), we compute the above 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 partial likelihood.

We use the Parametric Policy Iteration method (c.f. Benítez-Silva et al. (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. We defer the details of the numerical solution method to Appendix B.

Computing the log-likelihood function in (3) is difficult because it involves high dimensional integrations over the conditional distributions of serially correlated signals.[25] Because 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.[26] We defer the details of computing the likelihood to Appendix C.

Identification

The identification of the structural parameters in the pricing model is relatively straightforward. We do not estimate the daily discount factor δ, but calibrate it to match an

[25]Without conditioning on ξ_j, the signals y_i are correlated across periods.

[26]See Rubin (1988). See, also, Fernandez-Villaverde and Rubio-Ramirez (2007), Flury and Shephard (2008) and Gallant et al. (2009) for examples of applying SIR to simulate the likelihood function when estimating dynamic models. Alternatively, one can express $\tilde{l}\left(\left(p_{jt}, I_{jt}\right)_{t=1}^{T_j}|\theta_s\right)$ as follows:

$$\tilde{l}\left(\left(p_{jt}, I_{jt}\right)_{t=1}^{T_j}|\theta_s\right) = \int \tilde{l}\left(\left(p_{jt}, I_{jt}\right)_{t=1}^{T_j}|y^{T_j}, \theta_s\right) f\left(y^{T_j}\right) dy^{T_j},$$

where $f\left(y^{T_j}\right)$ is the probability density function of y^{T_j}. Then, one may compute $\tilde{l}\left(\left(p_{jt}, I_{jt}\right)_{t=1}^{T_j}|\theta_s\right)$ 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* in the estimation sample.

annual rate of 25%.[27] Among the structural parameters, ceteris paribus, a higher holding cost, ϕ_3, implies a lower initial price; a larger standard deviation of the signal from the initial assessment, σ_0, implies a smaller initial price dispersion and higher initial price. Meanwhile, fixing other parameters, larger σ_0 means greater residual uncertainty about ξ after the initial assessment. Therefore, the distribution of the initial prices and the total price changes due to the adjustments of the seller's belief about ξ help identify ϕ_3 and σ_0. The standard deviation of the signals received in the selling process, σ_s, determines the impact of these signals on the seller's belief about ξ. Changes in the seller's belief about ξ lead to price adjustments only when they are sufficient to justify the menu cost. Hence, the daily price-adjustment likelihood and the average magnitude of the observed one-time price adjustments help identify σ_s and ϕ_1 (the mean of the menu cost). Lastly, the effect of inventory on price, captured by parameter ϕ_2, is identified by the correlation between K_{jt} and p_{jt}.

5.3 Sample for Estimation

For the demand estimation, we use a subsample of the six most popular Japanese and Korean car models carried by the CarMax store in our data. These models include Honda Accord and Civic, Nissan Altima, Toyota Camry and Corolla, and Hyundai Sonata. We include only model years 2005-2010, dropping older cars. We focus on a small set of models with similar observable characteristics so that it is reasonable to assume that ξ is independent of car model and other observable attributes. We use the first day's data for each car in the first step of the estimation, and use the first six days' data in the second step. As discussed in the data section, we minimize the measurement error of sales by using only the first few days' data because cars may be delisted and taken to wholesale auctions after staying on the retail market for a long period of time.

There are 975 unique cars in the demand-estimation sample. Out of the 975 cars, 12.5% are sold on the first day, and 9.6% of the remaining cars are sold on the second day. Table 7 presents the summary statistics of these cars' attributes and first listing price. The listing prices on the first day range from \$8,599 to \$24,998, with the average being \$16,562. The mileage ranges from 3,100 miles to 117,250 miles, with an average of 32,530 miles. The dealer's daily inventory of these car models ranges from one to 39 cars.

For estimating the dynamic pricing model, we use the first six days' data for Honda

[27]The return to CarMax's invested capital (unleveraged) in 2011 was around 15% (c.f. CarMax's annual report 2011). On top of this, we assume an annual depreciation rate of 10% for cars.

Accords (the model with the largest number of cars in our data) from the above sample. Table 8 presents the summary statistics of this subsample of 178 unique cars. The prices and car attributes of this subsample are similar to those of the sample used for the demand estimation. As discussed earlier, the sample of the first few days is sufficient for identification. Further restriction to a single car model greatly reduces the computational cost of simulating the likelihood function of the price sequences, which increases linearly with the number of cars (and the number of days of each car) used in the estimation.[28] Because the observed pricing and sales patterns of CarMax do not vary much across car models (c.f. Table 6), we are likely to get similar estimates if using a subsample of other models.

6 Empirical Results

We first report the parameter estimates of the structural model and then show that our model is able to fit the pricing and sales patterns in the data reasonably well. Finally, we use the estimated structural model to quantify the values of the initial assessment and of the seller's subsequent learning in the selling process.

6.1 Model Estimates

Demand Model

Table 9 presents the estimated parameters in the reduced-form pricing equation. We include a full set of dummies for car model, model year and listing month in the pricing regression (and the demand model) to control for the fixed effects of these factors. The coefficients of the current inventories of the six models are all negative. These estimates are jointly significant at 5% confidence level and are consistent with the seller's incentive to avoid stock-out and that cars of the same model compete against each other. The estimated coefficients of other variables also seem reasonable. For example, the list prices are significantly higher for larger cars, cars with more powerful engines and those with lower mileage. It is worth noting that the adjusted R-squared of the regression is 0.74, meaning that there is still a fair amount of price variation that the variations in the observable attributes cannot explain.

Table 10 presents the parameters of the demand model, estimated using the adapted control function approach. The estimated price coefficient is -0.90 and is statistically

[28] It takes around eight days to estimate the dynamic pricing model using the subsample of Honda Accords.

significant at the 5% level. The estimated coefficients of other car attributes show that, for example, buyers' values for the cars increase with the engine volume (significant at the 10% level) and decrease with mileage (significant at the 5% level). The coefficient of price residual, ζ, is 0.69 (with the p-value being 10.9%), indicating a positive correlation between ξ and ζ.

The estimated standard deviation of η is 0.43, significant at the 1% level. Recall that the car's latent quality, ξ, can be expressed as $\psi\zeta + \eta$. Thus, we get the estimated standard deviation of ξ as $\sqrt{\hat{\psi}^2\hat{\sigma}_\zeta^2 + \hat{\sigma}_\eta^2}$, which is equal to 0.99. In comparison, the standard deviation of $X_j\beta$, the utility of observable attributes, is estimated to be 1.16. Thus, the across-car variation in latent car quality is similar in importance to that of the observable attributes.

For the choice of buying other cars of the same model, the coefficient of $log(K_{jt})$ is 3.95 (significant at the 10% level). This result shows that other cars of the same model compete for the same demand.

Dynamic Pricing Model

Table 11 presents the estimates of the structural parameters of the dynamic pricing model, which are all estimated accurately. The standard deviation of the signal from the initial assessment (σ_0) and that of the signal received in the selling process (σ_s) are estimated to be 0.48 and 0.64, respectively. The demand estimation shows that the seller's ex ante belief about ξ has a standard deviation of 0.99. These results suggest that (1) CarMax's initial assessment leads to much more accurate information about car-specific demand; and (2) the learning in the subsequent selling process is also very informative.

The estimated mean of the menu cost, ϕ_1, is \$41, which seems larger than the physical cost of changing prices. As discussed earlier, the menu cost may also include other costs, such as those for assessing new information and computing a new optimal price based on the updated belief. Furthermore, as we will explain in more detail later, the average menu cost actually incurred would be much smaller than the mean of the menu cost, because the seller can choose to change prices when the realized menu cost is relatively small.

The parameter that captures the effect of inventory on price, ϕ_2, is estimated to be -0.28. Thus, the seller has an incentive to set a lower price when the inventory is relatively high and to set a higher price when the inventory is relatively low. This is consistent with the seller's objective of avoiding stock-outs, as explained earlier. The holding cost is estimated to be \$71 per day, which seems reasonable given the maintenance cost and the opportunity cost of the parking space (when used to capacity).

Overall, our estimates seem to have good face validity. The following section shows that our estimated model fits the main pricing patterns in the data reasonably well.

6.2 Model Fit

To evaluate how well our model fits data, we simulate the price paths and sale outcomes for cars in the estimation sample for the pricing model. For each car, we first draw ξ_j from the distribution of $N(0, \sigma_\xi^2)$ and the initial assessment signal from the distribution of $N(\xi_j, \sigma_0^2)$. With the initial assessment signal, we update the seller's belief about ξ_j. The inventory on day one is taken directly from the data. Given these state variables, we compute the optimal initial price. If the car is not sold on day one, we draw a new signal from the distribution of $N(\xi_j, \sigma_s^2)$ and the inventory level for the next day; we then update the seller's belief about ξ_j, and compute the next day's optimal price. We simulate the buyer's purchase decision by using the estimated demand model. The process stops once the car is sold. We run the simulation 100 times for every car. The model predictions discussed below are the summaries of all the simulated data of all the cars in the estimation sample.

The top panel of Table 12 compares the predicted price levels and sale outcomes to those in the data. First, the average initial and transaction prices in the estimation sample are, respectively, \$17,315 and \$16,945, and those predicted by the model are, respectively, \$17,416 and \$17,211. The predicted average prices are close to those in the data, with the respective relative prediction errors being 0.6% and 1.6%. Second, in the estimation sample, 49.1%, 72.2% and 85.2% of the cars are sold within the first six, 15 and 30 days respectively; with the prices predicted by the pricing model, the demand model predicts that 58.9%, 88.1% and 98.4% of the cars are sold within the corresponding time frames. The predicted sales are all higher than the observed levels, with the respective relative prediction errors being 20%, 22% and 15%. A possible cause of the over-prediction of sales is that the actual demand's price elasticity is higher than what we estimated. The fit of sales seems reasonable given our relatively parsimonious demand model.

Overall, our model fits the important price and sales levels in the data reasonably well. These metrics are the key determinants of the seller's performance in the various scenarios of our policy experiments. Because the errors in these model predictions are relatively small, we expect their impact on the results of our policy experiments to be limited.

Recall that the most prominent feature of price dynamics in our data is the overall downward trend in the inter-temporal price movement. The bottom panel of Table 12 shows that, in the estimation sample, 23% of the cars experience total-price drops and 0.6% experience total-price increases,[29] while the corresponding model predictions are 32.5% and 15.0%. The average magnitudes of the total-price drops and increases are, respectively, \$597 and \$399 in the estimation sample, while the corresponding model predictions are \$486 and \$357.

Thus, our model is able to produce the overall downward trend in price adjustments, predicting many more price decreases than price increases. The average magnitudes of the predicted total price drop and increase are reasonably close to those in the data. The model, however, predicts more frequent price adjustments than appear in the data, and the over-prediction is much more significant for price increases. The significant over-prediction of the total-price increases seems a limitation of the Gaussian learning process that we assume in the pricing model. However, it is worth pointing out that, given the predictions of transaction prices and sales patterns, the impact of the prediction errors in the details of the price dynamics on our policy experiments would be small. This is because such prediction errors affect only the seller's average performance via the average total menu cost incurred in the selling process, which, as we will see below, is very small.

6.3 Policy Experiments

In this section, we conduct policy experiments to quantify the value of the information about car-specific demand that the seller obtains through the following two channels: the initial assessment and subsequent learning in the selling process. Our results shed light on the empirical importance of such information for dealers' profitability and for transaction efficiency in the used-car retail market.

For each (counterfactual) scenario, we simulate the price path and buyers' purchase decisions 100 times for every car in the sample for estimating the pricing model. There are four scenarios in the first set of experiments, defined by whether the initial assessment is conducted and the extent of the subsequent learning in the selling process: "Assessment and Learning," "Learning without Assessment," "Assessment and Weak Learning," and "Weak Learning without Assessment." Note that the scenario with "Assessment and Learning" is simply that of the estimated structural model. When there is no initial assessment, the seller's belief about ξ when setting the initial price

[29] We observe more (about 3%) cars with total price increases in the demand-estimation sample with all six models.

is just $N(0, \sigma_\xi^2)$. In the scenarios involving "weak learning," we set the standard deviation, σ_s, of the signals received in the selling process at 3 (the original estimate of σ_s is 0.64). To provide a benchmark, we also simulate the counterfactual scenario in which the seller has perfect information about ξ from the beginning.

We consider scenarios with weak learning instead of no learning at all because the former suits our empirical framework better and is informative enough for our objective. Particularly, for the scenario with neither initial assessment nor any subsequent learning, cars with low values of ξ would take a long time to sell and, thus, would eventually be taken off the retail channel and sold through wholesale auctions. Hence, it is difficult to evaluate the scenario without knowing the seller's expected wholesale prices. In contrast, with weak learning, cars are sold within a relatively short period; therefore, our simulation does not require the extra information on the expected wholesale prices.

Table 13 reports the average expected values of the following four metrics for the above five scenarios: the net revenue (net of the total holding cost and menu cost), the transaction price, the number of days until sale and the total menu cost. We define the "*value of the initial assessment*" as the difference in the average expected net revenue between "Assessment and Learning" and "Learning without Assessment," and the "*value of subsequent learning*" as that between "Learning without Assessment" and "Weak Learning without Assessment." This seems a natural way to define the two values because the decision of whether to conduct an assessment to improve the initial information is typically based on the premise that the seller already has achieved (at least to some extent) the benefit of subsequent learning in the selling process. Given our discussions in the previous paragraph, the "value of subsequent learning", as defined here, would be a conservative measure of the value of subsequent learning when there is no initial assessment.

The seller may achieve a higher expected profit by conducting an initial assessment because the information collected can lead to a) a higher expected transaction price; b) a lower total holding cost due to a shorter time to sell; and c) a smaller total menu cost due to less need to adjust prices. Our simulation results show that, relative to "Learning without Assessment," adding the initial assessment increases the average expected transaction price by \$65, lowers the total holding cost by \$21, and reduces the total menu cost by \$16. In total, the estimated value of the initial assessment is around \$101 per car.

The seller may benefit from subsequent learning in the selling process because it can lead to a higher expected transaction price and a lower total holding cost, at the expense of paying more menu costs. Our simulation results show that, relative to "Weak Learning without Assessment," adding learning increases the average transaction price by \$40 and lowers the total holding cost by \$239. These benefits are achieved at an additional menu cost of \$9. Overall, the estimated value of the subsequent learning is around \$269 per car (c.f. the summary in Table 14).

The above estimated values of information are significant, considering that the average net profit per car is about \$1150 for the cars in our estimation sample.[30] Therefore, these results show that subsequent learning in the selling process can significantly improve the profitability of dealers who conduct no initial assessment for their cars; even with the subsequent learning, a careful inspection of cars and research into local consumers' preferences still benefit the seller significantly.

It is worth pointing out that, in general, the initial assessment and the subsequent learning are substitutes. The results reported in Table 13 show that, when there is only weak learning, adding the initial assessment increases the net revenue by \$313, more than the value of the initial assessment estimated above. Similarly, when there is an initial assessment, changing weak learning to learning increases the net revenue by \$57, which is less than the value of the subsequent learning estimated above.

The benefits of the information-based pricing strategy must be weighed against the cost that dealers have to incur. The variable cost of thoroughly assessing each car should be lower than the estimated value of the initial assessment. The costs of developing and maintaining the necessary information and pricing system are likely to be large. Because of the substantial fixed costs, larger dealers would be in a better position to invest in and exploit such an information-based pricing strategy due to their greater economies of scale.

Table 13 also shows that, relative to "Assessment and Learning," the average expected net revenue can be further increased by \$129 if the seller begins with perfect information about car-specific demand. That improvement, compared to what has already been achieved, seems modest. This result is consistent with CarMax's alleged sophistication in collecting information and exploiting it in its pricing process.

Finally, we investigate whether menu cost prevents the seller from effectively utilizing the information learned in the selling process. Specifically, we simulate the

[30]The average transaction price in our estimation data is about \$17,200. The net profit rate per car is estimated based on the information in CarMax's annual report in 2011. The reported average selling price of used cars is \$18,019 (page 22). CarMax sold 396,181 used vehicles in 2011 (page 22). The total earning before taxes for used vehicles is about 492.8 million dollars, which is estimated based on information on page 41 and assumes the same earning for every revenue dollar. The net profit rate per used car is, then, 6.9%, calculated as the ratio of the net earning before tax per car to that of the average selling price.

counterfactual scenario of "Half Menu Cost," which modifies "Assessment and Learning" by lowering the mean of the menu cost by half. Table 15 shows that, with the menu cost cut by half, the average expected transaction price increases by $42; the average total menu cost incurred per car hardly changes; and the number of days until sale increases by 0.4 days. As a result, the average expected net revenue increases by only $19 per car.

The above results suggest that the menu cost (with the mean of $41) does not seem to significantly limit the seller's ability to utilize the information learned in the selling process. The intuition here is that, even though the menu cost is high on average, the seller can choose to change the price on days when the realized menu cost is low. Our results confirm that the average total menu cost actually incurred by the seller is small: in the case of "Assessment and Learning," for example, with 47% of cars experiencing at least one price adjustment, the average total menu cost incurred per car is only $4.8. Therefore, as long as the seller is not too impatient, the menu cost would not significantly impact her performance.

7 Summary and Concluding Remarks

The main challenge for dealers in pricing idiosyncratic products results from their ex ante lack of information about item-specific demand. In this paper, we developed a structural model of dynamic pricing for idiosyncratic products, featuring the optimal stopping structure and seller learning during the selling process. We show that seller learning impacts pricing dynamics through a rich set of mechanisms and tends to generate systematic price drops for each item over time. Our model-free analysis of sales data from the used-car retail market suggests that seller learning is a key factor in explaining the typically downward adjustments in the prices for individual cars.

We estimated the structural model of demand and dynamic pricing using a panel dataset of used-car sales from CarMax. The model fits the main patterns of price dynamics in the data. The policy experiments show significant value for the demand information that the dealer learns through initial assessment and in the subsequent selling process. These findings suggest a potentially high return to taking a more information-based approach to pricing idiosyncratic products.

Some aspects of our empirical framework may be improved or extended in future research. We adopted the Gaussian learning framework to describe the seller's learning process. Though it helps maintain the tractability of our model, it seems responsible for the significant over-prediction of upward price adjustments. A better model for the seller's learning process may help the model better fit the details of pricing dynamics.

We treated the seller's learning process for each item as independent of those for the other items. This is appropriate when $\xi_j \perp \xi_{j'} | (X_j, X_{j'})$; that is, the demands for individual items are independent of each other after conditioning on their observable attributes. Our empirical analysis uses data from a period in which the sales of used cars were quite stable, and, thus, the independence assumption seems reasonable for our case. However, when the market is fluctuating, the independence assumption can be too strong. In such situations, across-item learning may also be important: the information that the seller learns about the demand for one item may be also informative about the demand for another item. We leave it to future research to study the pricing for idiosyncratic products in such situations.

Tables and Figures

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

Note: The prices on the left-hand side are the Kelley Blue Book "Private Party" prices by car conditions in Rockville, Maryland, for the 2007 Honda Accord LX sedan with 68,500 miles.

Figure 2: Frequency of Price Adjustments by time on the market: CarMax

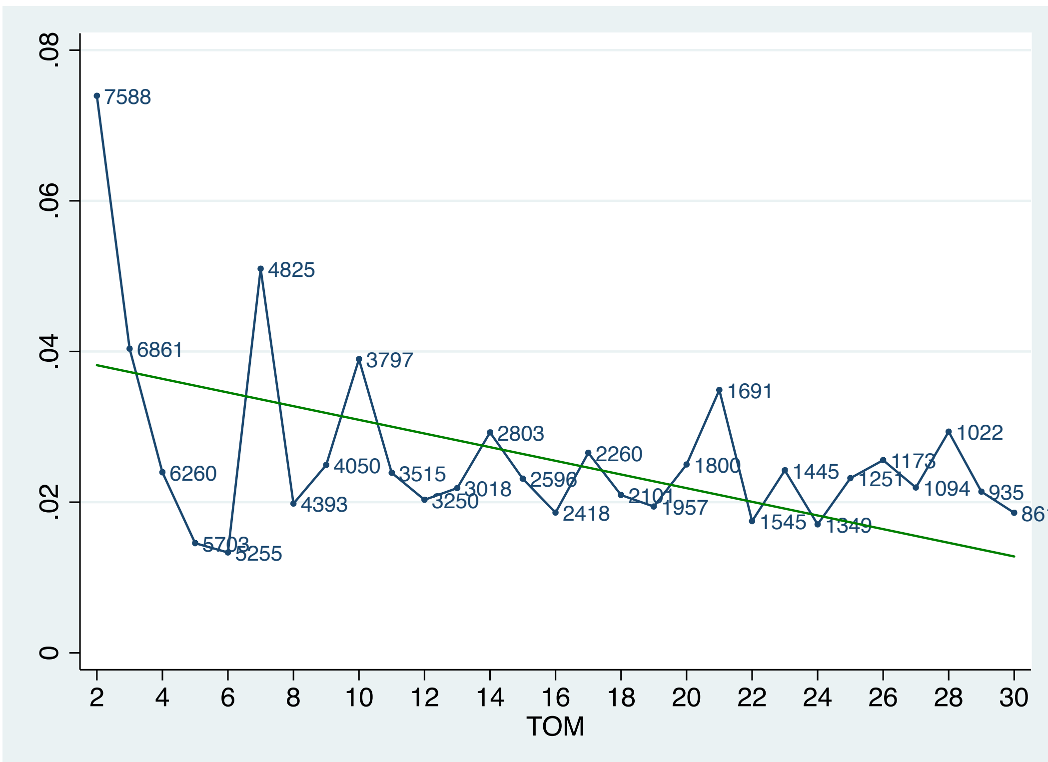

Note: For each day of time on the market, the dot represents the percentage of remaining cars with adjusted prices relative to the previous day, and the accompanying number is the number of cars that remain on the market on the particular day.

Figure 3: Frequency of Price Adjustments by time on the market: Other Dealers

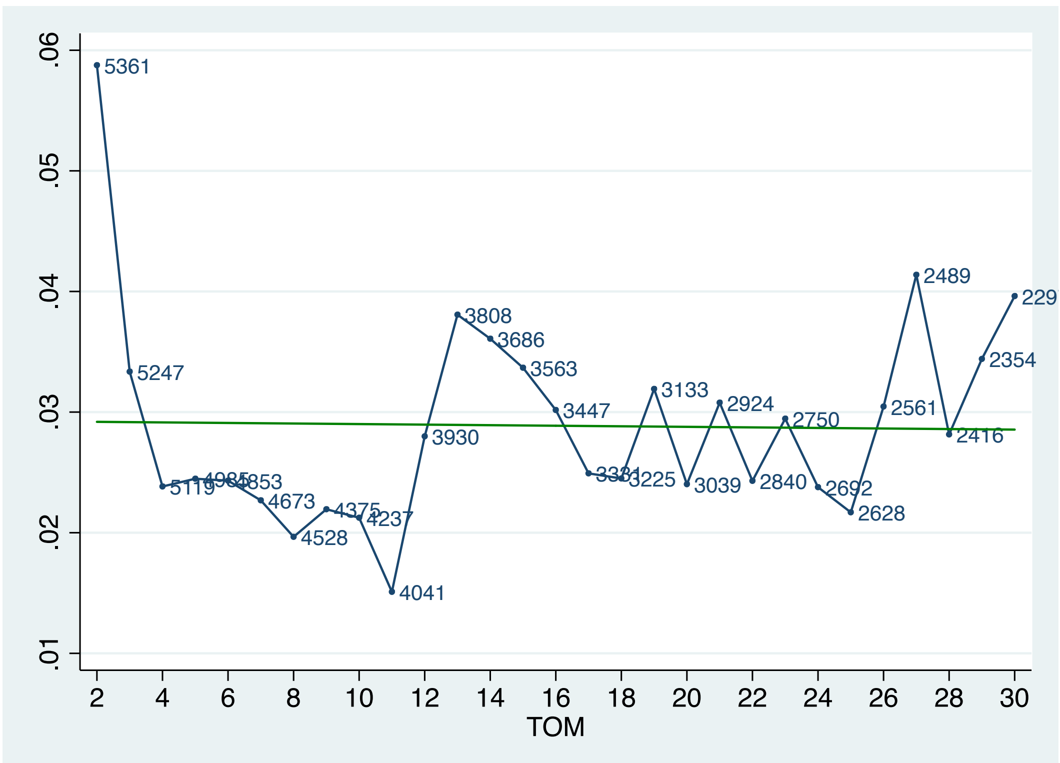

Note: This figure replicates Figure 2 using data from the five other largest local dealers (in terms of the number of unique cars listed in 2011). For each day of time on the market, the dot represents the percentage of remaining cars that have their prices adjusted relative to the previous day, and the accompanying number is the number of cars that remain on the market on the particular day.

Table 1: Distribution of Time to Sell

Dealer	N	Mean	Sd	Min	p25	p50	p75	Max
CarMax	7,630	14.0	13.8	1	4	9	19	120
Other large dealers	5,394	35.5	35.3	1	10	23	48	274

Note: A car's time to sell is defined as the number of days that the car remained on *Cars.com* until delisted. Other large dealers include five of the other largest local dealers (in terms of the number of unique cars listed in 2011).

Table 2: Total Number of Price Changes

	CarMax		Other large dealers	
Number of changes	Freq.	%	Freq.	%
0	5,350	70.1	2,926	54.3
1	1,754	23.0	1,116	20.7
2	402	5.3	602	11.2
3	96	1.3	284	5.3
4	18	0.2	210	3.9
5	7	0.1	109	2.0
6	1	0.01	63	1.2
7	2	0.03	43	0.8
8			23	0.4
9			11	0.2
10			3	0.06
11			4	0.07

Note: This table summarizes the total number of price changes during a car's entire duration on the market. Other large dealers include five of the other largest local dealers (in terms of the number of unique cars listed in 2011).

Table 3: Distribution of the Magnitudes of Price Changes

(a) CarMax

	N	Mean	Sd	Min	p25	p50	p75	Max
One-time price changes								
Price increases	257	731.9	385.6	10	399	601	1000	3000
Price decreases	2,715	499.9	345.7	3	399	399	601	3000
Total price changes								
Price increases	176	725.4	374.9	10	399	601	1000	2000
Price decreases	2,059	631.1	555.8	3	99	399	1000	6000

(b) Other large dealers

	N	Mean	Sd	Min	p25	p50	p75	Max
One-time price changes								
Price increases	688	1308.9	1485.4	1	504.5	1000	1492	16530
Price decreases	4,404	865.1	902.8	1	495	800	1000	16865
Total price changes								
Price increases	279	1258.9	1342.5	1	500	1000	1554	12210
Price decreases	2,141	1605.6	1356.3	1	824	1112	2018	16912

Note: This table summarizes the magnitudes of the price changes for all cars that experienced at least one price change. The price changes are in U.S. dollars. A one-time price change is defined as ("price on day t" - "price on day $t-1$"), and total price change is defined as ("price on the last day" - "price on the first day"). Other large dealers include five of the other largest local dealers (in terms of the number of unique cars listed in 2011).

Table 4: Daily Relative Change of Inventories

(a) CarMax

Variable	N	Mean	Sd	p10	p25	p50	p75	p90
All models	364	0.00	0.07	-0.08	-0.04	0.01	0.04	0.07
Top six car models								
Honda Accord	364	0.01	0.16	-0.14	-0.07	0	0.10	0.17
Nissan Altima	364	0.01	0.15	-0.17	-0.07	0	0.09	0.20
Toyota Camry	359	0.02	0.23	-0.20	0	0	0	0.20
Honda Civic	364	0.01	0.16	-0.17	0	0	0	0.20
Chevrolet Impala	347	0.02	0.25	-0.20	0	0	0	0.25
Chrysler Town & Country	338	0.03	0.29	-0.25	0	0	0	0.33

(b) Other large dealers

Dealer/No. 1 model	N	Mean	Sd	p10	p25	p50	p75	p90
Other dealer 1 / Camry	363	0.01	0.12	-0.10	0	0	0	0.13
Other dealer 2 / Cobalt	325	0.01	0.20	-0.11	0	0	0	0.00
Other dealer 3 / Silverado 1500	364	0.00	0.08	-0.09	0	0	0	0.00
Other dealer 4 / CTS	364	0.00	0.08	-0.07	0	0	0	0.09
Other dealer 5 / Jetta	364	0.01	0.11	-0.10	0	0	0	0.11

Note: The daily relative change of inventories for each model is defined as ("inventory on day t" - "inventory on day $t-1$")/"inventory on day $t-1$". Panel (a) summarizes the relative inventory change separately for the top six models for the CarMax store in our data; and Panel (b) summarizes the relative change in the inventory of the number one model at each of the other five largest dealers in the local market. Other large dealers include the five other largest local dealers (in terms of the number of unique cars listed in 2011).

Table 5: Time to Sell and Price Adjustments by Dealer

(a) Time to Sell

Variable	N	Mean	Sd	Min	p25	p50	p75	Max
CarMax	7630	14.0	13.8	1	4	9	19	120
Other dealer 1	1285	33.8	32.0	1	10	22	47	174
Other dealer 2	1261	23.7	18.1	1	9	18	36	85
Other dealer 3	1193	34.1	33.3	1	10	21	46	271
Other dealer 4	958	47.8	45.1	1	13	32	73	272
Other dealer 5	697	46.0	43.3	1	15	32	64	274

(b) Total Price Changes

Variable	N	Mean	Sd	Min	p25	p50	p75	Max
CarMax	7630	-153.6	427.3	-6000	-99	0	0	2000
Other dealer 1	1285	-390.3	1020.4	-10019	-500	0	0	3143
Other dealer 2	1261	-1001.0	1449.1	-14000	-2000	-600	0	6000
Other dealer 3	1193	-164.6	563.3	-4197	0	0	0	6510
Other dealer 4	957	-463.3	1236.0	-16912	-753	0	0	12210
Other dealer 5	690	-989.7	1789.1	-16865	-1983	-785	0	8010

Note: The time to sell of a car is defined as the number of days that the car remained on Cars.com until delisted. Total price change is defined as (price on the last day - price on the first day). The price changes are in U.S. dollars. Other dealers are the five other largest local dealers by the number of unique cars listed in 2011.

Table 6: Time to Sell and Price Adjustments by Model for CarMax

(a) Time to Sell

Variable	N	Mean	Sd	Min	p25	p50	p75	Max
Honda Accord	206	13.7	13.8	2	4	9	17	88
Nissan Altima	192	14.7	13.7	1	5	10	21	84
Toyota Camry	135	15.0	12.3	2	5	12	22	66
Honda Civic	163	16.7	16.7	1	4	11	22	92
Chevrolet Impala	116	13.8	12.1	1	5	10	19	81
Chrysler Town & Country	106	12.9	12.1	1	4	9.5	19	58

(b) Total Price Changes

Variable	N	Mean	Sd	Min	p25	p50	p75	Max
Honda Accord	206	-141.6	300.2	-1099	-99	0	0	1000
Nissan Altima	192	-174.8	358.7	-2000	-99	0	0	601
Toyota Camry	135	-137.7	325.3	-1399	0	0	0	601
Honda Civic	163	-101.8	321.9	-2000	0	0	0	1000
Chevrolet Impala	116	-87.0	328.5	-1399	0	0	0	2000
Chrysler Town & Country	106	-89.6	256.7	-1000	0	0	0	601

Note: The time to sell of a car is defined as the number of days that the car remained on Cars.com until delisted. Total price change is defined as (price on the last day - price on the first day). The price changes are in U.S. dollars.

Table 7: Estimation Sample (Demand Model) Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Price ($1,000)	16.56	2.55	8.60	24.00	975
Mile (10k)	3.25	1.93	0.31	11.73	975
Engine volume (liter)	2.44	0.50	1.7	3.5	975
Wheelbase	107.75	2.25	102	110	975
Model year 06	0.10	0.30	0	1	975
Model year 07	0.38	0.49	0	1	975
Model year 08	0.31	0.46	0	1	975
Model year 09	0.16	0.37	0	1	975
Altima	0.21	0.40	0	1	975
Camry	0.23	0.42	0	1	975
Civic	0.15	0.35	0	1	975
Corolla	0.09	0.29	0	1	975
Sonata	0.04	0.20	0	1	975
Seller's inventory of the same model	11.115	7.423	1	39	975

Note: The demand-estimation sample includes the top six models that the CarMax store in our data carried in 2011. These models are Honda Accord and Civic, Nissan Altima, Toyota Camry and Corolla, and Hyundai Sonata.

Table 8: Estimation Sample (Dynamic Pricing Model) Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Price ($1,000)	17.48	2.41	10.60	25.00	178
Mile (10k)	2.91	1.74	0.40	11.72	178
Engine volume (liter)	2.59	0.32	2.40	3.5	178
Wheelbase	108.57	0.91	108	110	178
Model year 06	0.10	0.29	0	1	178
Model year 07	0.60	0.49	0	1	178
Model year 08	0.17	0.38	0	1	178
Model year 09	0.12	0.32	0	1	178
Seller's inventory of the same model	19.13	8.51	3	39	178

Note: The estimation sample of the pricing model includes the Honda Accords the CarMax store in our data carried in 2011.

Table 9: Reduced-form Pricing Equation

Variable	Coefficients	Std. Err.
Constant	-14.17	8.08
Mile (10k miles)	-0.56***	0.03
Engine volume (liter)	2.30***	0.13
Wheelbase	0.21***	0.07
Inventory	-0.004	0.02
Inventory*Altima	-0.08	0.06
Inventory*Camry	-0.07**	0.03
Inventory*Civic	-0.03	0.04
Inventory*Corolla	-0.10	0.12
Inventory*Sonata	-0.45	0.33
Dummies for car model, model year, month	Yes	
Num. of Observations	975	
Adj. R^2	0.743	

Note: This table reports the regression results of the reduced-form pricing equation using the first day's data of each car in the demand-estimation sample. The dependent variable is the first day's list price measured in \$1,000. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 10: Multinomial Probit Demand Model Parameter Estimates

Variables	Coefficients	Std. Err.
Constant	-7.51	10.83
Price ($1,000)	-0.90**	0.44
Mile (10k miles)	-0.50**	0.25
Engine volume (liter)	1.88*	1.02
Wheelbase	0.17	0.13
Price residual (ζ)	0.69	0.43
Std. dev. of η:	0.43***	0.07
Choosing other cars of the same model		
Constant	-9.95	5.90
Log(# of other cars)	3.95*	2.31
Bivariate normal distribution of $(\varepsilon_{jt}, \varepsilon_{-jt})$		
Covariance of $(\varepsilon_{jt}, \varepsilon_{-jt})$	-3.74	2.48
Std. dev. of ε_{-jt}	5.79*	3.20
Dummies for car model, model year, time	Yes	
Num. of cars	975	

Note: The estimation sample include the first six days' data (all the data if sold before the sixth days) of each car in the demand-estimation sample. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 11: Dynamic Pricing Model Parameter Estimates

Parameters	Coefficients	Std. Err.
Std. dev. of the initial assessment signal		
σ_0	0.48***	0.03
Std. dev. of signals received after the initial assessment		
σ_s	0.64***	0.08
Mean of the menu cost		
ϕ_1	$40.79***	11.09
The inventory effect		
ϕ_2	-0.28***	0.09
Holding cost		
ϕ_3	$70.67***	7.41
Num. of cars	178	

Note: The dynamic pricing model is estimated using the first six days' data (all the data if sold before the sixth days) of each Honda Accord that the CarMax store in our sample carried in 2011. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 12: Model Fit

Performance metrics	Data	Model prediction
Price and sales levels		
Average initial price	17,315	17,416
Average transaction price	16,945	17,211
Cars sold in the first six days (%)	49.1	58.9
Cars sold in the first 15 days (%)	72.2	88.1
Cars sold in the first 30 days (%)	85.2	98.4
Inter-temporal price adjustments		
Cars with total price decrease (%)	22.5	32.5
Average total price decrease	597	486
Cars with total price increase (%)	0.6	15.0
Average total price increase	399	357

Note: The reported outcome metrics are the averages over all the simulated data for cars in the estimation sample for the pricing model. These metrics can be interpreted as the average expected outcomes across different cars because they are the same as first taking the average over the 100 simulations for each car and then taking the average over these averaged numbers. The prices are in U.S. dollars.

Table 13: Policy Experiment Results

Scenarios	Net revenue	Transaction price	Time to Sell	Total menu cost
Assessment and Learning	16,680	17,211	7.4	4.8
Learning w/o Assessment	16,579	17,146	7.7	20.4
Assessment and Weak Learning	16,623	17,272	9.2	1.1
Weak Learning w/o Assessment	16,310	17,106	11.1	14.1
Perfect Initial Information	16,809	17,353	7.7	0.0

Note: For each scenario, the reported outcome metrics are the averages over all the simulated data for cars in the estimation sample for the pricing model. These metrics can be interpreted as the average expected outcomes across different cars because they are the same as first taking the average over the 100 simulations for each car and then taking the average over these averaged numbers. The net revenue, transaction price and total menu cost are all measured in U.S. dollars.

Table 14: The Values of the Initial Assessment and Subsequent Learning

(Assessment, Learning)	*Yes*	*Weak*	Value of subsequent learning
Yes	16,680	16,623	
No	16,579	16,310	269
Value of initial assessment	101		

Note: The values reported for the four scenarios are the average expected net revenues taken from Table 13.

Table 15: The Impact of Menu Cost

Scenarios	Net revenue	Transaction price	Time to Sell	Total menu cost
Assessment and Learning	16,751	17,211	7.4	4.8
Half Menu Cost	16,770	17,253	7.8	4.6

Note: The scenario "Assessment and Learning" is the same as that in Table 13. The simulated scenario "Half Menu Cost" is obtained in a similar way, except that the mean menu cost is cut by half. The net revenue, transaction price and total menu cost are all measured in U.S. dollars.

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. Technical report, 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. Technical report, Yale University.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica 63*(4), 841–890.

Blackwell, D. (1953). Equivalent comparison of experiments. *Annals of Mathematical Statistics 24*, 265–272.

Ching, A. (2010). A dynamic oligopoly structural model for the prescription drug market after patent expiration. *International Economic Review 51*, 1175–1207.

Ching, A., T. Erdem, and M. Keane (2011). Learning models: An assessment of progress, challenges and new developments. Technical report, 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. Technical report, Oxford University.

Gallant, R., H. Hong, and A. Khwaja (2009). Estimating a dynamic oligopolistic game with serially correlated unobserved production costs. Technical report, Duke University.

Grossman, S. J., R. E. Kihlstrom, and L. J. Mirman (1977). A bayesian approach to the production of information and learning by doing. *Review of Economic Studies 44*, 533–547.

Larsen, B. (2014). The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. Technical report, National Bureau of Economic Research.

Lewis, G. (2011). Asymmetric information, adverse selection and online disclosure: The case of ebay motors. *The American Economic Review*, 1535–1546.

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.

Merlo, A. and F. Ortalo-Magne (2004). Bargaining over residential real estate: evidence from england. *Journal of Urban Economics 56*(2), 192–216.

Merlo, A., F. Ortalo-Magne, and J. Rust (2013). The home selling problem: Theory and evidence.

Mirman, L., L. Samuelson, and A. Urbano (1993). Monopoly experimentation. *Internatiaonl 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.

Riley, J. and R. Zeckhauser (1983). Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics 98*(2), 267.

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 J. Bernardo, M. DeGroot, D. Lindley, and A. Smith (Eds.), *Bayesian Statistics 3*. Oxford University Press.

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of harold zurcher. *Econometrica 55*, 999–1033.

Sweeting, A. (2012). Dynamic pricing behavior in perishable goods markets: Evidence from secondary markets for major league baseball tickets. *Journal of Political Economy 120*(6), 1133–1172.

Tadelis, S. and F. Zettelmeyer (2014). Information disclosure as a matching mechanism: Theory and evidence from a field experiment.

Taylor, C. (1999). Time-on-the-market as a sign of quality. *Review of Economic Studies 66*(3), 555–578.

Trefler, D. (1993). The ignorant monopolist: Optimal learning with endogenous information. *International Economic Review 34*, 565–581.

Xu, X. and W. Hopp (2005). Dynamic pricing and inventory control: The value of demand learning. Technical report, Northwestern University.

Appendix A: Transforming the Seller's Profit Maximization Problem into a Sequential Optimization Problem

In this appendix, we show that the seller's original profit maximization problem can be transformed into a sequential optimization problem. First, note that the profit-maximization problem can be reformulated as follows:

$$\begin{aligned}
&\max_{(p_t)_{t=1}^{\infty}} \left\{ E_{\xi} E_{(y^t)_{t=1}^{\infty}|\xi} \left(E_{(K_t)_{t=2}^{\infty}|K_1} \sum_{t=1}^{\infty} \delta^{t-1} \chi_t E_{\varphi_t} \pi_t \left(p_t, \xi, K_t, \varphi_t\right) \right) \right\} \\
&= E_{y^1} \max_{(p_t)_{t=1}^{\infty}} \left\{ E_{\xi|y^1} E_{(y^t)_{t=2}^{\infty}|\xi, y^1} \left(E_{(K_t)_{t=2}^{\infty}|K_1} \sum_{t=1}^{\infty} \delta^{t-1} \chi_t E_{\varphi_t} \pi_t \left(p_t, \xi, K_t, \varphi_t\right) \right) \right\},
\end{aligned} \tag{4}$$

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^t that the seller has received by the beginning of day t, we have:

$$\begin{aligned}
&\max_{(p_\tau)_{\tau=t}^{\infty}} \left\{ E_{\xi|y^t} E_{(y^\tau)_{\tau=t+1}^{\infty}|\xi, y^t} \left(E_{(K_\tau)_{\tau=t+1}^{\infty}|K_t} \sum_{\tau=t}^{\infty} \delta^{\tau-t} \chi_\tau E_{\varphi_\tau} \pi_\tau \left(p_\tau, \xi, K_\tau, \varphi_\tau\right) \right) \right\} \\
&= \max_{p_t} \left\{ E_{\xi|y^t} E_{\varphi_t} \pi_t \left(p_t, \xi, K_t, \varphi_t\right) + \max_{(p_\tau)_{\tau=t+1}^{\infty}} E_{\xi|y^t} \left\{ \left(1 - D\left(p_t, K_t, \xi\right)\right) \right. \right. \\
&\quad \left. \left. \delta E_{(y^\tau)_{\tau=t+1}^{\infty}|\xi, y^t} E_{(K_\tau)_{\tau=t+1}^{\infty}|K_t} \sum_{\tau=t+1}^{\infty} \delta^{\tau-(t+1)} \chi_\tau E_{\varphi_\tau} \pi_\tau \left(p_\tau, \xi, K_\tau, \varphi_\tau\right) \right\} \right\} \\
&= \max_{p_t} \left\{ E_{\xi|y^t} E_{\varphi_t} \pi_t \left(p_t, \xi, K_t, \varphi_t\right) + E_{\xi|y^t} \left(1 - D\left(p_t, K_t, \xi\right)\right) E_{y^{t+1}|\xi, y^t} E_{K_{t+1}|K_t} \right. \\
&\quad \left. \delta \max_{(p_\tau)_{\tau=t+1}^{\infty}} E_{\xi|y^{t+1}} \left\{ E_{(y^\tau)_{\tau=t+2}^{\infty}|\xi, y^{t+1}} E_{(K_\tau)_{\tau=t+2}^{\infty}|K_{t+1}} \sum_{\tau=t+1}^{\infty} \delta^{\tau-(t+1)} \chi_\tau E_{\varphi_\tau} \pi_\tau \left(p_\tau, \xi, K_\tau, \varphi_\tau\right) \right\} \right\},
\end{aligned} \tag{5}$$

where the second equality follows by the Law of Iterated Expectations. Taken together, reformulations (4) and (5) imply that the seller's profit optimization problem can be transformed into a sequential optimization problem, which has a Bellman Equation representation.

Appendix B: Numerical Solution of the Dynamic Pricing Model

In this appendix, we describe in detail the numerical method that we use to solve the dynamic pricing model presented in the model section (Section 4). The notations are the same as in Section 4. Our objective is to solve the following Bellman equation:

$$V(S_t) = E_{\varphi_t} \max_{p_t} \left\{ E_{\xi|y^t} E_{\varepsilon_t} \pi_t (p_t, \xi, K_t, \varphi_t) + E_{\xi|y^t} (1 - D(p_t, K_t, \xi)) \delta E_{y^{t+1}|\xi, y^t} E_{K_{t+1}|K_t} V(S_{t+1}) \right\}$$

$$\begin{aligned} s.t.\ \pi_t &= -\varphi_t \cdot 1\{p_t \neq p_{t-1}\} + (p_t - m(K_t)) D(p_t, K_t, \xi) - \phi_3 \\ m(K_t) &= \left(\frac{2\exp\left(\phi_2\left(K_t - \bar{K}\right)\right)}{1 + \exp\left(\phi_2\left(K_t - \bar{K}\right)\right)} - 1 \right) \bar{c} \\ \mu_{t+1} &= \frac{\sigma_t^2 y_t + \sigma_s^2 \mu_t}{\sigma_t^2 + \sigma_s^2} \\ \sigma_{t+1}^2 &= \frac{\sigma_t^2 \sigma_s^2}{\sigma_t^2 + \sigma_s^2} \\ y_t &= \xi + \epsilon_t, \end{aligned}$$

where $S_t \equiv ((\mu(y^t), \sigma_t), p_{t-1}, K_t)$. Among the state variables, $(\mu(y^t), \sigma_t, p_{t-1})$ are continuous variables; K_t is a discrete state variables; φ_t is a random variable of the exponential distribution with mean ϕ_1. We use V^* to denote the unique solution to the above Bellman equation.

We use the Parametric Policy Iteration algorithm to numerically solve the above Bellman equation (as an analytical solution is not available). In the algorithm, 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:

$$\begin{aligned} & V(\mu_t, \sigma_t, p_{t-1}, K_t) \\ = & \sum_{l=1}^{L} \Psi_l \rho_l (\mu_t, \sigma_t, p_{t-1}, K_t) \\ \equiv & \rho(\mu_t, \sigma_t, p_{t-1}, K_t) \Psi, \end{aligned}$$

where $\rho_l(\mu_t, \sigma_t, p_{t-1}, K_t)$ are multivariate basis functions, $\rho \equiv (\rho_1, \rho_2, ..., \rho_L)$ and $\Psi \equiv (\Psi_1, \Psi_2, ..., \Psi_L)'$. In our application, we use the Chebyshev polynomials as the basis functions. Let $D(p) \equiv D(p, K_t, \xi)$, and define $\bar{p}(S_t|V)$ as follows:

$$\bar{p}(S_t|V) = \arg\max_p E_{\xi|y^t} \left\{ (p - m(K_t)) D_t(p) + (1 - D_t(p)) \delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} V(S_{t+1}) \right\}.$$

That is, $\bar{p}(S_t|V^*)$ is the optimal pricing strategy if the seller decides to update the price. And let $W(S_t|V)$ be the value function defined as follows:

$$W(S_t|V) = \max_p E_{\xi|y^t} \left\{ (p - m(K_t)) D_t(p) + (1 - D_t(p)) \delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} V(S_{t+1}) \right\},$$

and $U(S_t|V)$ be defined as follows:

$$U(S_t|V) = E_{\xi|y^t} \left\{ (p_{t-1} - m(K_t)) D_t(p_{t-1}) + (1 - D_t(p_{t-1})) \delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} V(S_{t+1}) \right\}$$

Thus, $W(S_t|V^*)$ is the seller's value function, ignoring the current day's menu cost, and $U(S_t|V^*)$ is the seller's expected payoff if she chooses not to change the price. It is obvious that we have $W(S_t|V) \geq U(S_t|V)$, and that:

$$V^*(S_t) = E_{\varphi_t} \max \left\{ W(S_t|V^*) - \varphi_t, U(S_t|V^*) \right\}.$$

Define $\bar{\varphi}(S_t|V) \equiv W(S_t|V) - U(S_t|V)$. Then, the seller updates the price if and only if $\varphi_t < \bar{\varphi}(S_t|V^*)$. Let the distribution function of φ_t be F. Then, we have:

$$V^*(S_t) = F(\bar{\varphi}(S_t|V^*)) W(S_t|V^*) + (1 - F(\bar{\varphi}(S_t|V^*))) U(S_t|V^*) - \int_0^{\bar{\varphi}(S_t|V^*)} \varphi_t dF(\varphi_t). \tag{6}$$

The optimal pricing strategy can be represented as follows:

$$p^*(S_t) = \begin{cases} \bar{p}(S_t|V^*) & \text{, if } \varphi_t < \bar{\varphi}(S_t|V^*) \\ p_{t-1} & \text{, otherwise,} \end{cases}$$

which is completely determined by $(\bar{p}(S_t|V^*), \bar{\varphi}(S_t|V^*))$.

The Parametric Policy Iteration algorithm starts with an initial guess of the optimal pricing strategy, $(\bar{p}_0(S_t), \bar{\varphi}_0(S_t))$, and involves two steps for each iteration. Let subscript k be the iteration number. Then, the first step is the policy-evaluation step: updating the value function given the pricing strategy $(\bar{p}_{k-1}(S_t), \bar{\varphi}_{k-1}(S_t))$ resulting from the last iteration (we suppress the argument of the threshold function $\bar{\varphi}_{k-1}(S_t)$ in the following to simplify notation). In this step, we solve the following linear functional equation for $V_k(S_t)$:

$$\begin{aligned}V_k(S_t) =& F(\bar{\varphi}_{k-1}) E_{\xi|y^t}\Big\{(p_{k-1}(S_t) - m(K_t)) D_t(p_{k-1}(S_t)) + \\ & (1 - D_t(p_{k-1}(S_t))) \delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} V_k(S_{t+1}(p_{k-1}(S_t)))\Big\} + \\ & (1 - F(\bar{\varphi}_{k-1})) E_{\xi|y^t}\Big\{(p_{t-1} - m(K_t)) D_t(p_{t-1}) + \\ & (1 - D_t(p_{t-1})) \delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} V_k(S_{t+1}(p_{t-1}))\Big\} - \int_0^{\bar{\varphi}_{k-1}} \varphi_t dF(\varphi_t),\end{aligned}$$

where $S_{t+1}(p) \equiv ((\mu(y^{t+1}), \sigma_{t+1}), p, K_{t+1})$.[31] The above equation becomes the following system of linear equations in Ψ after substituting in the approximating polynomial for $V_k(S_t)$:

$$\begin{aligned}\rho(S_t)\Psi = & F(\bar{\varphi}_{k-1}) E_{\xi|y^t}(p_{k-1}(S_t) - m(K_t)) D_t(p_{k-1}(S_t)) + \\ & F(\bar{\varphi}_{k-1}) E_{\xi|y^t}(1 - D_t(p_{k-1}(S_t))) \delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} \rho(S_{t+1}(p_{k-1}(S_t)))\Psi \\ & (1 - F(\bar{\varphi}_{k-1})) E_{\xi|y^t}(p_{t-1} - m(K_t)) D_t(p_{t-1}) + \\ & (1 - F(\bar{\varphi}_{k-1})) E_{\xi|y^t}(1 - D_t(p_{t-1})) \delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} \rho(S_{t+1}(p_{t-1}))\Psi - \\ & \int_0^{\bar{\varphi}_{k-1}} \varphi_t dF(\varphi_t).\end{aligned}$$

After combining terms with the same coefficients, we have:

$$\varrho_k(S_t)\Psi = Y_k(S_t)$$

where:

$$\begin{aligned}\varrho_k(S_t) \equiv & \rho(S_t) - F(\bar{\varphi}_{k-1}) E_{\xi|y^t}(1 - D_t(p_{k-1}(S_t))) \delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} \rho(S_{t+1}(p_{k-1}(S_t))) - \\ & (1 - F(\bar{\varphi}_{k-1})) E_{\xi|y^t}(1 - D_t(p_{t-1})) \delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} \rho(S_{t+1}(p_{t-1})) \\ Y_k(S_t) \equiv & F(\bar{\varphi}_{k-1}) E_{\xi|y^t}(p_{k-1}(S_t) - m(K_t)) D_t(p_{k-1}(S_t)) + \\ & (1 - F(\bar{\varphi}_{k-1})) E_{\xi|y^t}(p_{t-1} - m(K_t)) D_t(p_{t-1}) - \int_0^{\bar{\varphi}_{k-1}} \varphi_t dF(\varphi_t).\end{aligned}$$

If we pick a grid of N different points of S_t: $(S_{1,}..., S_N)$, then we can solve for Ψ by

[31] This equation is an analog of equation (6), with $(V^*, \bar{p}(S_t|V^*), \bar{\varphi}(S_t|V^*))$ replaced by $(V_k, \bar{p}_{k-1}(S_t), \bar{\varphi}_{k-1})$

using the least squares criterion as follows:

$$\Psi = \arg\min_{\tilde{\Psi}} \sum_{i=1}^{N} \left(\varrho_k \left(S_i \right) \tilde{\Psi} - Y_k \left(S_i \right) \right)^2 .$$

The grid points we use are the so-called Chebyshev points, which result in minimum approximation error and avoid the "Runge's phenomenon" associated with uniform grid points.

Given $V_k \left(S_t \right) = \rho \left(S_t \right) \Psi$, we can carry out the second step—the policy function improvement step—by solving the optimal prices at the grid points as follows:

$$\begin{aligned} \bar{p}_k \left(S_t \right) &= \bar{p} \left(S_t | V_k \right) \\ \bar{\varphi}_k \left(S_t \right) &= W \left(S_t | V_k \right) - U \left(S_t | V_k \right) . \end{aligned}$$

To implement the algorithm, we start with an initial guess of $\left(\bar{p}_0 \left(S_t \right), \bar{\varphi}_0 \left(S_t \right) \right)$, and then 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 value function in the Bellman equation. The optimal pricing policy can be easily computed with the solution of the value function.

Appendix C: Simulating the Partial Likelihood

In this appendix, we describe in detail how we compute the partial likelihood in (3). We apply the Sampling and Importance Resampling method to simulate the integrations in the partial likelihood function. For preparation of the following discussion, define $\wp_1(\mu_t, \sigma_t, K_t, p_{t-1}, \varphi_t)$ as the optimal pricing strategy given the old price p_{t-1}, and define $\wp_0(\mu_t, \sigma_t, K_t)$ as the optimal pricing strategy when ignoring the current menu cost. That is:

$$
\begin{aligned}
&\wp_1(\mu_t, \sigma_t, K_t, p_{t-1}, \varphi_t) \\
&= \arg\max_p \varphi_t 1\{p \neq p_{t-1}\} + \\
&E_{\xi|y^t}\left\{(p - m(K_t)) D_t(p) + (1 - D_t(p)) \delta E_{y^{t+1}|\xi, y^t} E_{K_{t+1}|K_t} V(S_{t+1})\right\} \\
&\wp_0(\mu_t, \sigma_t, K_t) \\
&= \arg\max_p E_{\xi|y^t}\left\{(p - m(K_t)) D_t(p) + (1 - D_t(p)) \delta E_{y^{t+1}|\xi, y^t} E_{K_{t+1}|K_t} V(S_{t+1})\right\}.
\end{aligned}
$$

It is clear that $\wp_1(\mu_t, \sigma_t, K_t, p_{t-1}, \varphi_t) = \wp_0(\mu_t, \sigma_t, K_t)$ if $\wp_1(\mu_t, \sigma_t, K_t, p_{t-1}, \varphi_t) \neq p_{t-1}$, and $\wp_0(\mu_1, \sigma_1, K_1)$ is the optimal price on the first day. Furthermore, let $\bar{\varphi}(\mu_t, \sigma_t, K_t, p_{t-1})$ be the menu cost that makes the seller indifferent between charging $\wp_0(\mu_t, \sigma_t, K_t)$ and keeping the old price p_{t-1}. So, we have:

$$
\wp(\mu_t, \sigma_t, K_t, p_{t-1}, \varphi_t) = \begin{cases} \wp_0(\mu_t, \sigma_t, K_t) & \text{if } \varphi_t \leqslant \bar{\varphi}(\mu_t, \sigma_t, K_t, p_{t-1}) \\ p_{t-1} & \text{otherwise.} \end{cases}
$$

Recall that in our model, we have that $y_{jt} \equiv \xi_j + \epsilon_{jt}$ ($\xi_j \perp \epsilon_{jt}$ for all t), where $\xi_j \sim N(0, \sigma_\xi^2)$, $\epsilon_{j0} \sim N(\xi_j, \sigma_0^2)$ and $\epsilon_{jt} \sim N(\xi_j, \sigma_s^2)$ for $t \geq 1$. Thus, we have that $f(\xi_j|y_{j0}) = \frac{1}{\sigma_1}\phi\left(\frac{\xi_j - \mu_{j1}}{\sigma_1}\right)$ (ϕ is the density function of the standard normal distribution), where $\mu_{j1} = \frac{\sigma_\xi^2 y_{j0}}{\sigma_\xi^2 + \sigma_0^2}$ and $\sigma_1^2 = \frac{\sigma_0^2 \sigma_\xi^2}{\sigma_\xi^2 + \sigma_0^2}$, and $f(\xi_j|y_j^t) = \frac{1}{\sigma_t}\phi\left(\frac{\xi_j - \mu_{jt}}{\sigma_t}\right)$, where $y_j^t \equiv (y_{j0}, ..., y_{j,t-1})$, $\mu_{jt} = \frac{\sigma_{t-1}^2 y_{j,t-1} + \sigma_s^2 \mu_{t-1}}{\sigma_{t-1}^2 + \sigma_s^2}$ and $\sigma_t^2 = \frac{\sigma_{t-1}^2 \sigma_s^2}{\sigma_{t-1}^2 + \sigma_s^2}$. In addition, it is easy to verify that $f(y_{jt}|y_j^t) = f(y_{jt}|\mu_{jt}) = \frac{1}{\sqrt{\sigma_t^2 + \sigma_s}}\phi\left(\frac{y_{jt} - \mu_{jt}}{\sqrt{\sigma_t^2 + \sigma_s}}\right)$ and $f(\mu_{t+1}|\mu_t) = \frac{\sqrt{\sigma_t^2 + \sigma_s^2}}{\sigma_t^2}\phi\left(\frac{\mu_{t+1} - \mu_t}{\sigma_t^2 / \sqrt{\sigma_t^2 + \sigma_s^2}}\right)$.

The partial likelihood we need to compute is:

$$
\tilde{l}\left((p_{jt}, I_{jt})_{t=1}^{T_j}\right) = l(p_{j1}) \prod_{t=2}^{T_j} \int l\left(p_{jt}|y_j^t, p_{j,t-1}\right) f\left(y_j^t|(p_{j\tau}, I_{j\tau})_{\tau=1}^{t-1}\right) dy_j^t.
$$

First, we have that:

$$l\left(p_{j1}\right) \quad = \quad f_{\mu_{j1}}\left(\mu_1\left(p_{j1}\right)\right) / \left|\frac{\partial \wp_0\left(\mu_1\left(p_{j1}\right), \sigma_1, K_{j1}\right)}{\partial \mu_1}\right|,$$

where $\mu_t\left(p_{jt}\right)$ is defined by $p_{jt} = \wp_0\left(\mu_t, \sigma_t, K_{jt}\right)$ and $f_{\mu_{j1}}\left(\mu_1\right) = \frac{1}{\frac{\sigma_\xi^2}{\sqrt{\sigma_\xi^2+\sigma_0^2}}}\phi\left(\frac{\mu_1}{\frac{\sigma_\xi^2}{\sqrt{\sigma_\xi^2+\sigma_0^2}}}\right)$. We compute the partial derivative $\frac{\partial \wp_0(\mu_1(p_{j1}),\sigma_1,K_{j1})}{\partial \mu_1}$ by applying the implicit function theorem to the first-order condition of the Bellman equation.[32] Note that

$$\int l\left(p_{jt}|y_j^t, p_{j,t-1}\right) f\left(y_j^t|\left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-1}\right) dy_j^t = \int l\left(p_{jt}|\mu_{jt}, \sigma_t, p_{j,t-1}\right) f\left(\mu_{jt}|\left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-1}\right) d\mu_{jt}.$$

Given the ns simulated random draws of $\left(y_{j,s}^t\right)_{s=1}^{ns}$ from the conditional distribution of $f\left(y_j^t|\left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-1}\right)$, we can get the corresponding random sample of $\left(\mu_{jt,s}\right)_{s=1}^{ns}$ from the conditional distribution of $f\left(\mu_{jt}|\left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-1}\right)$. There are two different cases for computing $\int l\left(p_{jt}|\mu_{jt}, \sigma_t, p_{j,t-1}\right) f\left(\mu_{jt}|\left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-1}\right) d\mu_{jt}$. In discussing the first case, we use $p_t\left(\mu_{jt}\right)$ to stand for $\wp_0\left(\mu_{jt}, \sigma_t, K_t\right)$ in some places to simplify notation, and we use $\delta_{\tilde{x}}\left(x\right)$ to denote the Dirac delta function of x that has point mass of one at $\tilde{x}$ and is zero elsewhere. For the first case of $p_{jt} \neq p_{j,t-1}$, we have

$$\begin{aligned}
& \int l\left(p_{jt}|\mu_{jt}, p_{j,t-1}\right) f\left(\mu_{jt}|\left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-1}\right) d\mu_{jt} \\
= & \int \delta_{p_{jt}}\left(p_t\left(\mu_{jt}\right)\right) \Pr\left(p_{jt} \neq p_{j,t-1}|\mu_{jt}\right) f\left(\mu_{jt}|\left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-1}\right) d\mu_{jt} \\
= & \int \frac{\delta_{\mu_t(p_{jt})}\left(\mu_{jt}\right) \Pr\left(p_{jt} \neq p_{j,t-1}|\mu_{jt}\right)}{\left|\frac{\partial \wp_0(\mu_t(p_{jt}),\sigma_t,K_{jt})}{\partial \mu_t}\right|} f\left(\mu_{jt}|\left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-1}\right) d\mu_{jt} \\
= & f\left(\mu_t\left(p_{jt}\right)|\left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-1}\right) \Pr\left(p_{jt} \neq p_{j,t-1}|\mu_t\left(p_{jt}\right)\right) / \left|\frac{\partial \wp_0\left(\mu_t\left(p_{jt}\right), \sigma_t, K_{jt}\right)}{\partial \mu_t}\right|
\end{aligned}$$

where $f\left(\mu_t\left(p_{jt}\right)|\left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-1}\right)$ can be computed using Bayes' rule as follows:

$$\begin{aligned}
& f\left(\mu_t|\left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-1}\right) \\
= & \frac{\Pr\left(I_{j,t-1} = 0|p_{j,t-1}, \left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-2}, \mu_t\right) f\left(\mu_t|p_{j,t-1}, \left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-2}\right)}{\Pr\left(I_{j,t-1} = 0|p_{j,t-1}, \left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-2}\right)}
\end{aligned}$$

[32]Note that we can let μ and y incorporate $X_j\beta$, so that we do not need to solve the dynamic pricing problem separately for each value of $X_j\beta$. The observation greatly reduces the computational burden in our estimation.

We can get the two probabilities in the above expression by simulation. For the density of $f\left(\mu_t|p_{j,t-1},(p_{j\tau},I_{j\tau})_{\tau=1}^{t-2}\right)$ in the above equation, we have that:

$$\begin{aligned} & f\left(\mu_t|p_{j,t-1},(p_{j\tau},I_{j\tau})_{\tau=1}^{t-2}\right) \\ = & \int f\left(\mu_t|\mu_{t-1}\right) f\left(\mu_{t-1}|p_{j,t-1},(p_{j\tau},I_{j\tau})_{\tau=1}^{t-2}\right) d\mu_{t-1}, \end{aligned}$$

which equals $f\left(\mu_t|\mu_{t-1}\left(p_{j,t-1}\right)\right)$ for the special case of $p_{j,t-1}\neq p_{j,t-2}$. For the second case, $p_{jt}=p_{j,t-1}$, we have:

$$\begin{aligned} & \int l\left(p_{jt}|\mu_{jt},p_{j,t-1}\right) f\left(\mu_{jt}|\left(p_{j\tau},I_{j\tau}\right)_{\tau=1}^{t-1}\right) d\mu_{jt} \\ = & \frac{1}{ns}\sum_{s=1}^{ns}\Pr\left(p_{jt}=p_{j,t-1}|\mu_{jt,s}\right) f\left(\mu_{jt,s}|\left(p_{j\tau},I_{j\tau}\right)_{\tau=1}^{t-1}\right). \end{aligned}$$

We use the Sampling and Importance Resampling method to generate random samples from the distribution of $f\left(y_j^t|\left(p_{j\tau},I_{j\tau}\right)_{\tau=1}^{t-1}\right)$ and $f\left(y_j^{t+1}|\left(p_{j\tau},I_{j\tau}\right)_{\tau=1}^{t-1},p_{jt}\right)$. For the purpose of simulating the integrations, we use the corresponding random sample of $(\mu_{jt})_{s=1}^{ns}$. It is sufficient for us to keep the record of the simulated random sample in the form of $(y_{jt,s},\mu_{jt,s})$ instead of the entire vector of signals received. When $p_{jt}\neq p_{j,t-1}$, we just update the simulation sample with ns draws of the same value $\mu_{jt}=\mu_t\left(p_{jt}\right)$. The actual simulation procedure proceeds in the following steps:

1. First, compute the value of $\mu_1\left(p_{j1}\right)$ as the "random sample" from the $f\left(\mu_{j1}|p_{j1}\right)$; then, simulate ns random draws from $f\left(y_{j1}|p_{j1}\right)$, by drawing $\left(y_{j1,s}\right)_{s=1}^{ns}$ from the distribution of $f\left(y_{j1}|\mu_1\left(p_{j1}\right)\right)$.

2. Filter step: resampling from $\left(y_{j1,s},\mu_1\left(p_{j1}\right)\right)_{s=1}^{ns}$ by using $\Pr\left(I_{j1}=0|\left(y_{j1,s},\mu_1\left(p_{j1}\right)\right)\right)$ as the sampling weight, which produces a random sample from $f\left(y_j^2|\left(I_{j1},p_{j1}\right)\right)$. From it, we can get the corresponding random sample of $f\left(\mu_{j2}|\left(I_{j1},p_{j1}\right)\right)$, which is used to simulate the integration in computing the likelihood of $l\left(p_{j2}|\left(I_{j1},p_{j1}\right)\right)$.

3. Filtering step: if $p_{j2}=p_{j1}$, then, resample from the last random sample by using $\Pr\left(p_{j2}=p_{j1}|\mu_{j2,s},p_{j1}\right)$ as the sampling weight; otherwise, replace the entire sample with ns repeated values of $\mu_2\left(p_{j2}\right)$. Thus, it generates the random sample of $f\left(\mu_{j2}|\left(I_{j1},p_{j1}\right),p_{j2}\right)$.

4. Prediction step: for each $\mu_{j2,s}$, draw a $y_{j2,s}$ from the distribution of $f\left(y_{j2}|\mu_{j2,s}\right)$, which produces a random sample of $\left(y_{j2,s},\mu_{j2,s}\right)_{s=1}^{ns}$ from the conditional distribution of $f\left(\left(y_{j2},\mu_{j2,s}\right)|\left(I_{j1},p_{j1}\right),p_{j2}\right)$

Repeating steps 3 and 4, we generate all the random samples that we need for simulating the integrations.