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

Guofang Huang

Hong Luo

Jing Xia

May, 2018

Abstract

Pricing products such as used cars, houses, and artwork is often challenging, because each item is unique, and the seller, ex ante, lacks information about the demand for individual items. This paper develops a dynamic-pricing model for products with significant item-specific demand uncertainty, in which a forward-looking seller learns about the item-specific demand through an initial assessment, as well as during the selling process. The model demonstrates how seller, through several mechanisms, learning can lead to the commonly observed downward trend in the prices of individual items. These mechanisms include the seller's optimal adjustment of prices over time to account for the dynamic adverse selection of unsold items and the diminishing option value in future learning. The model is estimated using novel panel data of a leading used-car dealership. Counterfactual experiments show that the value of learning in the selling process is \$203 per car. Conditional on subsequent learning in the selling process, the initial assessment further improves profit per car by \$139. With the dealer's net profit per car being about \$1150, these estimates suggest a potentially high return to taking an 'information-based' approach towards pricing products with item-specific demand uncertainty.

Keywords: dynamic pricing, item-specific demand, demand uncertainty, supply-side learning.

---

*Guofang Huang: gfhuang@andrew.cmu.edu; Hong Luo: hluo@hbs.edu; Jing Xia: jingxia@fas.harvard.edu. We thank Dirk Bergemann, Guido Imbens, Przemek Jeziorski, Zhenyu Lai, Ariel Pakes, Greg Lewis, Jiwoong Shin, K. Sudhir, Juuso Valimaki and Zhixiang Zhang for helpful discussions. Any remaining errors are our own.

1 Introduction

This paper studies the pricing problem for dealers selling products such as used cars, houses, and artwork. A defining feature of these products is that they exhibit significant *item-specific* heterogeneity even after accounting for all of their standard observable attributes.[1] Take the example of used cars. Identical new cars can end up as used cars in very different conditions after logging the same mileage, depending on how they were driven and maintained. For example, 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 "Fair condition" to \$12,550 for "Excellent condition."[2]

The main challenge in pricing this type of product is the difficulty of assessing demand for individual items. Dealers, *a priori*, often lack information on the latent quality/condition of specific items because they are, by definition, not the prior owners and often acquire such products in large numbers from the wholesale market with limited information on each individual item. Furthermore, the demand for a specific item also depends on local buyers' preferences, which may change over time. To address this challenge, dealers may acquire more information on the demand for each item through multiple channels. They may inspect and research the demand for the items they have acquired before setting the initial price. They also can learn more information in the subsequent selling process by, for example, observing instances in which the item does not sell and by 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.[3]

[1] The particular type of products that this paper focuses on can also be characterized by its "single-unit availability," as there are rarely two identical used cars (houses, paintings) available. These products are also similar to the so-called "idiosyncratic products" studied in Einav et al. (2018).

[2] 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.

[3] Conversations with a former top executive and a former employee working in CarMax's pricing department provided us an overview of the company's pricing system. Their descriptions are also consistent with CarMax's account in its annual reports (e.g., Annual Report 2011, pages 5 and 8). CarMax thoroughly inspects every car, checking over 125 points, before putting it up for sale. The inspection covers the engine and all major systems, including "cooling, fuel, drivetrain, transmission, electronics, brakes, steering, air conditioning and other equipment, as well as the interior and exterior of the vehicle." It maintains a proprietary database of 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 has sat on the lot, and when test drives occur. The company's proprietary pricing system uses all of this information, both from the initial inspections and information obtained from the sales process (such as data tracked by the RFID tags), to generate pricing recommendations. With these recommendations from the centralized pricing department, the managers at local stores have the final say over setting prices and making adjustments. Austin Ligon, the former CEO of CarMax, described its proprietary information and pricing systems as "one of our biggest competitive advantages"

Though potentially beneficial, the information-based pricing approach requires costly investments (e.g., in the information and pricing systems) and nontrivial fixed and variable costs. Should firms make these investments? More specifically, what is the value of the information that sellers can acquire through their initial assessment and subsequent learning in the selling process? From a theoretical perspective, how does the seller's learning in the selling process affect the optimal pricing strategy? To answer these questions, we develop a structural model of dynamic pricing with seller learning. We estimate the model using a novel panel dataset from a CarMax store, and conduct counterfactual experiments to quantify the value of the demand information that the dealer acquires before and during the selling process.

Our theoretical model of dynamic pricing involves a forward-looking seller selling a single item to sequentially arriving buyers. The seller is uncertain about an item-specific demand parameter. Before setting the initial price, the seller receives an unbiased signal—a signal that quantifies the result of the seller's initial assessment—about the demand parameter. In the subsequent selling process, the seller receives a new signal about the demand parameter 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 use the Bayesian Gaussian learning framework to formalize the seller's learning process. The seller incurs a cost whenever she adjusts the price, and her objective is to set prices based on her latest information to maximize the present value of her expected profit from selling the item. In essence, what the seller in our model faces is an optimal-stopping problem, in which the prices that she sets control the probability of a sale (stopping).

The above model captures several mechanisms through which seller learning leads to the stylized downward trend in the prices of individual items. First, the dynamic adverse selection of unsold items causes the optimal price to go down over time. The fact that the current buyer does not purchase an item implies that the seller is likely to have overestimated its demand. Thus, the seller is more likely to adjust her belief and, hence, the price downward after learning new information about demand. Apart from the dynamic adverse-selection effect, 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 in order to delay sale and, hence, benefit from the

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).

new information. Because the value of information decreases as the seller becomes more informed, the incentive to delay sale weakens over time. As a result, the optimal price can decrease even when the seller's expectation of the item-specific demand parameter remains constant. 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 possible value of the item-specific demand parameter, 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 (which we will later refer to as the 'active-learning' incentive) is stronger at the beginning of the sales process, leading to additional price drops under certain conditions.

We estimate our structural model using a car-level panel dataset from a CarMax store. The data include detailed car attributes and daily list prices (which are non-negotiable) over each car's duration on the market. We estimate the demand model using the control function approach (c.f. Petrin and Train (2010)) and the structural dynamic-pricing model 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 the item-specific demand parameter 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 random variables that are serially correlated, we compute the observable likelihood by simulation, using the method of Sampling and Importance Re-sampling.

Our policy experiments show that the total value of information acquired on car-specific demand —through both the initial assessment and learning during the sales process—is \$342. Given that the dealer's net profit per car is about \$1150 for our estimation sample, the value of information is substantial. Thus, the information-based pricing approach is worth considering, especially for large dealers that can implement it at scale. Furthermore, our results show that even without the initial assessment, the subsequent learning alone can improve the seller's profit by \$203. This suggests that improving the seller's ability to learn from new information revealed through the selling process is a useful starting point if substantial upfront investment is required to obtain a precise initial assessment.

This paper contributes to the literature on dynamic pricing by providing a novel empirical study of the dynamic-pricing problem for products with substantial item-specific demand uncertainty.

First, the paper's empirical framework makes a methodological contribution. It highlights such essential features as the optimal-stopping structure of this particular type of pricing problem and the effect of dynamic adverse selection on price dynamics. With some modifications, the framework can be applied to the pricing problems of other products with the same characteristics. Second, the paper makes some new theoretical observations about the optimal pricing strategy in the presence of seller learning. In particular, it shows that even if the seller's expectation of the demand does not change, the optimal price can still decrease as the residual uncertainty (and, hence, the option value of drawing more signals) decreases over time. Third, we empirically measure the value of an information-based pricing approach, providing useful references for firms in a market of significant economic importance.[4]

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3 introduces the data and presents model-free evidence of the uncertainty dealers face 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, and Section 7 concludes.

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 new products. We focus on the pricing problem for products exhibiting significant item-specific heterogeneity, which is distinguished by its inherent *optimal-stopping* structure. Our structural model is most closely related to that in Mason and Välimäki (2011), which studies the general problem of optimal stopping when the environment changes because of learning. Our model differs from theirs

[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.

[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 approaches different from ours. We are also not going into the large literature on revenue management for the same reason.

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, which is modeled as a continuous variable, derives from the seller's lack of information about item-specific heterogeneity or the preference of local consumers. Because of this difference, our model is able to capture additional pricing dynamics such as those driven by the dynamics in the value of new information.[6] In addition, our model may also better capture pricing problems in which the seller's information about the condition of and local preference for individual items is limited—such as problems faced by used-car dealerships or by banks selling a large number of foreclosed houses,.

Riley and Zeckhauser (1983) is another closely related theoretical paper that studies the optimal mechanism for selling an item to sequentially arriving buyers. The authors 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 we empirically quantify the value of information about item-specific demand. Also related is Handel and Misra (2015), which studies the optimal dynamic pricing strategies for non-storable new products. In their model, the seller knows the set of possible demand functions but does not have a prior belief about the likelihood of each. The key difference between Handel and Misra (2015) and this paper is that theirs focuses on optimal approaches to pricing new products with unlimited inventory, whereas ours focuses on empirically quantifying the value of information in pricing products with single-unit availability.

The empirical literature on dynamic pricing for products with single-unit availability is limted. A related paper from this literature is Merlo et al. (2015). Their paper builds a structural empirical model that helps explain a number of stylized facts about the pricing dynamics of individual houses, which Merlo and Ortalo-Magne (2004) also documents, using rich panel data from the UK. Our

[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, interdependent: 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.

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.

The broader empirical literature on dynamic pricing focuses mainly on pricing problems for new products without supply constraint. In a paper closely related to ours, Ching (2010b) 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' dynamic 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. By using the estimated and calibrated parameter values to simulate his dynamic oligopoly structural model, Ching finds that his model is able to match the observed increase of brand-name drugs' prices and the observed decrease of generic drugs' prices after generic entry. The main difference between our paper and Ching (2010b) is that: (i) we also deal with the dynamic adverse selection of unsold items and the role of inventory, in addition to dynamic pricing and seller learning; (ii) however, we consider dynamic pricing only by a monopoly, while Ching (2010b) explicitly models dynamic pricing in an oligopolistic environment. Another related paper, Nair (2007) shows that consumers' forward-looking behavior significantly limits the seller's ability to price discriminate intertemporally. Also related is Newberry (2016), which shows that a "demand-based" pricing scheme employed by an online music market raises consumer surplus but lowers the online market's expected revenue.

Our paper also relates closely to Hitsch (2006), which studies how the opportunity to learn sequentially about the demand for new products through sales data affects a firm's decisions about product launch, advertising, and product termination. Hitsch's dynamic model of product launch and exit highlights the importance of accounting for the option value of delaying scraping a product (to gain more information). His application of the model to the data of the ready-to-eat cereal industry shows that 1) the value of reducing demand uncertainty can be substantial; 2) firms may want to increase the rate of product launch when demand uncertainty is high; and 3) firms' uncertainty about demand for new products can explain the observed high launch and exit rates. We also focus on the general question of the impact of demand uncertainty and supply-side learning on firms' marketing-mix decisions, even though our specific research questions are very different.

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 online retail auctions mitigates the classic adverse-selection problem. Through a large-scale field experiment, Tadelis and Zettelmeyer (2015) 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 bargaining due to asymmetric information in this market is small. In contrast to these, we focus on dynamic pricing as a selling mechanism used by dealers in the offline used-car retail market.

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).[7] Our focus on the supply-side problem of pricing products with significant item-specific demand uncertainty sets our paper apart from these papers.

[7]Ching et al. (2013) provide a comprehensive survey of the empirical literature on consumer learning, and Ching et al. (2017) survey more recent developments in learning dynamics.

3 Data and Model-Free Analysis

3.1 Data

The data used in this paper are obtained 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 2010 and 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 the listing dealers sell them,[8] the data also allow us to determine the dates on which cars were sold. It is worth noting that the dealers in our sample listed their inventory on *Cars.com* only for advertising purposes: all cars were sold through the dealers' brick-and-mortar stores.

We focus on CarMax in our empirical analysis for the following reasons. First, as mentioned above, CarMax provides 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 listed its entire inventory on *Cars.com* during our data period.[9] Accurate information on cars' transaction prices and brand-level inventories are important for analyzing the dealer's pricing behavior. Finally, CarMax stores' inventory is often many times larger than that of their largest competitors in their local markets. The firm's 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 removed from the website, possibly because they were taken to wholesale auctions instead of being sold in the retail channel. These tended to be cars that had been on the market for a long time. In the case of CarMax, this issue should cause little concern since, as CarMax's 2011 Annual Report explained: "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, to estimate our structural

[8] According to people in 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 cars on the internet 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.

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

model, it is sufficient to use only the first few days' data for each car; 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 to explain. To reduce potential variance driven by time trend, our model-free analysis in the following uses only data in 2011. The 2010 data produce very similar patterns.

3.2 Model-Free Analysis

Stylized Pricing Patterns

We start with CarMax's pricing patterns. First, cars usually take a few days to sell, giving the seller enough opportunities 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% of its cars are sold within 31 days.

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 increases and decreases are $1308.9 and $865.1, respectively; 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 the sample period and that most cars were sold within a relatively short time, the price changes in our data are most likely driven by two main 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 to 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 CarMax's total inventory, as well as 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 Section 4.3, 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.

Alternative Hypotheses

In a standard setting of dynamic pricing for *new* products, heterogeneity in buyers' valuations can lead to downward price dynamics (price skimming) when the seller practices inter-temporal price discrimination (e.g., Nair (2007)). In particular, the seller may set a relatively high initial price to first sell to buyers with high valuations. Then, the seller would lower the price to sell to buyers with lower valuations. The price-skimming hypothesis, however, does not directly explain the observed downward price trend in our case. In contrast to new products, of which there are many identical units available for sale, each used car is unique. The single-unit availability implies that the seller cannot sell the same used car to a high-value buyer first and to a low-value buyer later.

That said, however, there is an interesting connection between the price-skimming strategy in a standard setting and our hypothesis of dynamic pricing with seller learning, because we can think of the latter as a type of "virtual" price skimming. With standard price skimming, the seller clears the market of consumers with high values by actually selling the products to them; in our setting, however, as a used car stays on the market, the seller gradually clears the possible high values of its latent quality out of her belief.

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 possible reason that a car has not been sold is that the previous buyers have unfavorable information about its quality. Thus, 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 the price would decrease with time on the market.

The available empirical evidence, however, does not support the signaling effect as a major driving force of the price patterns in our data. Although consumers can filter cars listed on *Cars.com* by their time on the market (up to a few ranges of days on the market), a sample of consumer browsing data that we obtained from the website shows that less than 1% of the searches used that option. We also investigate this possible explanation empirically by including in the demand model[10] a set of dummies that indicate the number of days that a car has been on the market. The estimated coefficients of all the dummies are statistically insignificant (results not reported in the paper), suggesting that time on the market has limited effect on demand in our setting.

Overall, the two alternative theories discussed above seem unlikely to be the major drivers of the pricing dynamics that we have documented. That said, it is important to note that we cannot completely rule them out as potential factors influencing price changes.

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 less time to sell its cars (c.f. Table 5a) and adjusts prices for a smaller share of its cars (c.f. Table 2). CarMax's total price adjustments are, on average, smaller than those made by other dealers (c.f. Table 5b). Furthermore, CarMax's performance is consistent across its available brands (c.f. Table 6). 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.[11]

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 other products of the same type, such as houses and artwork.

[10]The demand model is introduced later in Section 4.2.

[11]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.

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 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 are 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 products such as used cars, conditional on observable characteristics, the seller's uncertainty and learning about demand is mostly item-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, which is observable to all buyers but not to the researchers.

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 not buying car j be $v_{0t} = \varepsilon_{0t}$. Our specification is motivated partly by the need to keep the demand-side model parsimonious so that the corresponding dynamic pricing model is computationally tractable.[12]

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

$$I_{jt} = 1\{v_{jt} > v_{0t}\}.$$

That is, buyer t buys car j if and only if the choice gives him a higher value than the outside option.

Lastly, we assume that 1) ξ_j is independent of X_j, and 2) for all t, the preference shocks $\varepsilon_t \equiv (\varepsilon_{jt}, \varepsilon_{0t})$ are independent random draws from the same bivariate normal distribution and are independent of (X_j, 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. In practice, the uncertainty about the demand for a used car comes from two main sources: 1) the car's idiosyncratic quality; and 2) local buyers' preference for the car, given its standard attributes and its idiosyncratic quality. Both types of uncertainty can be substantial before the seller takes costly actions to mitigate them, but after a thorough inspection of the car's

[12] A previous version of the paper allowed for the option of buying cars other than car j in j's segment, and the expected payoff of the option was approximated by an affine function of $log(K_{jt})$ (where K_{jt} is the number of other cars in j's segment on the same day). The current, simpler specification, combined with the dynamic pricing model, better fits the dynamic price and sales patterns. This is likely because, with a simpler specification, the demand-model parameters are more precisely estimated.

condition, the seller's residual uncertainty should mostly be about local buyers' preference for a particular used car.[13]

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. The assessment can include: 1) visual and mechanical inspections of the car; and 2) research on local consumers' preference for the car based on historical data. 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. The seller can learn by observing that no sale takes place or by directly communicating with the 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.

It is important to note that there are alternative ways to model seller learning in the selling process. In particular, the Bernoulli-learning framework assumes that ξ is a Bernoulli random variable and that the seller learns by observing that a car has not been sold by the end of each day (see, for example, Zhang (2010) and Mason and Välimäki (2011) for applications of the alternative framework). An appealing feature of this approach is that it allows the price to directly affect the informativeness of the no-sale events. For example, if the price is set very high, a no-sale observation should reveal little information. The Bernoulli-learning framework also implies somewhat different learning dynamics. In particular, it implies that the seller's expectation of ξ decreases deterministically over time. In contrast, the Gaussian-learning model predicts that the seller's expectation of ξ is more likely to decrease over time, without ruling out the possibility of it increasing occasionally.[14]

[13] Note that our model allows for the possibility that a dealer, without conducting any inspection after acquiring the car from the wholesale auction, could have worse information about a car's mechanical quality than a buyer who has carefully test driven the car (or taken it to an independent mechanic for inspection). Because we have no direct evidence of how thorough CarMax's inspection is (except for the information self-reported by CarMax), we empirically estimate the informativeness of the initial assessment from the dynamic pricing patterns. The idea is that the more thorough the initial inspection is, the less residual uncertainty will remain, and, thus, we should observe a smaller adjustment in the total price.

[14] Specifically, the Bernoulli-learning model implies that the seller's expectation of ξ, $E_t(\xi)$, decreases deterministically over time, while the variance of her belief, $var_t(\xi) = E_t(\xi) * (1 - E_t(\xi))$, can either decrease deterministically

We choose the Gaussian-learning framework because the large, continuous support of Gaussian random variables fits our application quite well. First, CarMax's mechanical inspection covers over 125 points for each car, implying at least 2^{125} possible values for the initial-inspection signal. Thus, a continuous random variable is likely to provide a better approximation of the inspection signal and the subsequent beliefs.[15] Second, the large continuous support of the Gaussian distribution, in principal, allows our model to rationalize all observed prices. In contrast, with the Bernoulli-learning framework, the two mass points of ξ estimated from the demand side limit the range of prices that can be rationalized by the pricing model, leading to zero likelihood for some price observations in the estimation.

With the Gaussian-learning framework, one may alternatively 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 no-sale instances (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} determines not only the amount of information the seller learns each day, but also 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.

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

(if $E_t(\xi) <= 1/2$ for $t = 1$) or first increase and then decrease (if $E_t(\xi) > 1/2$ for $t = 1$). In contrast, the Gaussian learning model implies that the seller's expectation of ξ is more likely to decrease over time without ruling out the possibility of it going up from time to time (when the seller receives a sufficiently positive signal), while the variance of the seller's belief decreases deterministically. In practice, the seller may receive credible positive signals for unsold cars in certain scenarios. For example, a buyer may have a hard time choosing between two cars. In such situations, the salesperson may receive positive comments from the buyer about the car that she eventually decides not to buy and should interpret such information as a positive signal.

[15] Suppose that the outcomes from the large number of different inspection points are i.i.d. Bernoulli random variables, the result from the initial inspection would be a binomial random variable that can also be approximated by a Gaussian random variable.

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.[16] 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. We define K_{jt} as the number of cars, *excluding* car j, in car j's segment on day t, and specify the opportunity cost as $m(K_{jt}; \phi_2)$. 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. Thus, the seller may want to set prices lower to sell cars 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.[17] Empirically, we define the segment as cars of the same model as car j. As explained in Section 5.3, we use cars of model year between 2005 and 2009 in the analysis. Thus, a car's segment excludes cars older than six years, and, implicitly, cars with very high mileage.[18]

We assume that K_{jt}, conditional on $I_{jt} = 0$ (i.e., car j *not* having been sold by the end of day t), evolves as an exogenous first-order Markov process. The transition probability matrix $\Pr(K_{j(t+1)}|(K_{jt}, I_{jt} = 0))$ is determined by the probability of other cars in car j's segment being

[16] 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 were assumed to be fixed, and functions with kinks would be difficult to approximate accurately.

[17] See, also, Ishihara and Ching (2016) for a discussion of a similar effect of inventory on used-good pricing.

[18] Additional reduced-form pricing equations that include other inventory variables show that cars of different models or cars of the same model but of vintages older than 2005 have much smaller impacts on the focal car's price. Furthermore, the inventory of cars of the same model but of different model years (within the 2005-2009 range) has a significant effect on prices. Thus, the above definition of segment seems reasonable.

sold on day t (given $I_{jt} = 0$), as well as by the arrival process of new inventory in car j's segment. For our application, conditioned on $I_{jt} = 0$, the assumption of an exogenous stochastic process of K_{jt} (and, thus, independent of car j's price) seems reasonable for the following reasons. First, even though we cannot prove it for Probit demand models, for multinomial Logit models, the property that the probability of other cars in car j's segment being sold on day t given $I_{jt} = 0$ (and, thus, $\Pr(K_{j(t+1)}|(K_{jt}, I_{jt} = 0)))$ is independent of car j's price (and attributes) is implied by the independence-of-irrelevant alternatives (IIA) property of Logit models.

Second, empirically, as shown in Table 4, the inventory variable K_t shows no trend in either direction (that is, the frequency and magnitudes of inventory increases and decreases are mostly symmetric), even though prices mainly decrease over time. To further check the assumption, Table 7 presents the regression of K_{jt} on $K_{j(t-1)}$ and $p_{j(t-1)}$ using our demand-estimation sample. The regression incorporates the condition $I_{j(t-1)} = 0$ automatically, because our sample includes only observations of cars until they are sold, and, thus, observing K_{jt} implies that $I_{j(t-1)} = 0$. The regressions show that the coefficients of p_{t-1} are statistically insignificant in both OLS and fixed-effect regressions, and adding p_{t-1} has a negligible impact on the R-squared of the regressions.

We also 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.

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 the car 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 treat the demand as constant for each car's entire duration on the market, which seems a reasonable approximation for our empirical application. We use the function $D(p_{jt}, \xi_j) \equiv E_{\varepsilon_t} I_{jt}(X_j, p_{jt}, \xi_j, \varepsilon_t)$

to denote buyer t's probability of buying car j conditional on (p_{jt}, ξ_j), emphasizing its dependence on price p_{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, $\left(y^t, K_t, p_{t-1}, \varphi_t\right)$, 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.[19] Formally, the seller's profit-maximization problem can be written as follows:

$$
\begin{aligned}
& \max_{(p_t)_{t=1}^{\infty}} E_{\xi} E_{(y^t)_{t=1}^{\infty}|\xi} E_{(K_t)_{t=2}^{\infty}|K_1} \sum_{t=1}^{\infty} \delta^{t-1} \chi_t E_{\varphi_t} \pi_t \left(p_t, \xi, K_t, \varphi_t\right), \\
\text{s.t. } \pi_t &= \left(p_t - m\left(K_t\right)\right) D\left(p_t, \xi\right) - \phi_3, \text{ if } t = 1, \\
\pi_t &= -\varphi_t \cdot 1\left\{p_t \neq p_{t-1}\right\} + \left(p_t - m\left(K_t\right)\right) D\left(p_t, \xi\right) - \phi_3, \text{ if } t \geq 2, \\
m\left(K_t\right) &= \left(\frac{2\exp\left(\phi_2\left(K_t - \bar{K}\right)\right)}{1 + \exp\left(\phi_2\left(K_t - \bar{K}\right)\right)} - 1\right)\bar{c}, \\
\chi_t &= \prod_{\tau=1}^{t-1}\left(1 - I_\tau\right), \\
y^{t+1} &= \left(y^t, y_t\right), \\
y_t &= \xi + \epsilon_t,
\end{aligned}
$$

where χ_t indicates the availability of the car at the beginning of day t; δ is the seller's discount factor; $m\left(K_t\right)$ 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 $m\left(K_t\right)$ 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.

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:

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

[19] 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.

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

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

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

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:

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

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

where $S_t \equiv \left(\left(\mu\left(y^t\right), \sigma_t\right), p_{t-1}, K_t\right)$, and the value function depends on y^t only through $\left(\mu\left(y^t\right), \sigma_t^2\right)$, 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 $1 - D(p_t, \xi)$, with which the car stays unsold and the seller receives the corresponding option value of selling to future buyers, $\delta E_{y^{t+1}|\xi,y^t} E_{K_{t+1}|K_t} V(S_{t+1})$.[20] 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

In the model, seller 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 the 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 overestimates the 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.[21] Therefore, with this 'dynamic adverse selection' of unsold cars, learning creates a downward trend in the optimal prices.

Second, the dynamics in the value of new information generate additional intertemporal decreases 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 of new information leads to higher option values and higher optimal current prices (relative to the scenario without new information available on day t but is, otherwise, the same as in our model).[22] Because the value of new information

[20] 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.

[21] 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 overestimated 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.

[22] Suppressing the adverse-selection effect, we can express the value of new information as

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.[23] 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 $\frac{\partial^2 D(p_t,\xi)}{\partial p_t \partial \xi} < 0$ for all ξ. Additional dynamics in the optimal price arise as these incentives weaken over time, but the exact price dynamics that the active-learning effect creates depend on the shape of the demand function $D(p_t, \xi)$ and how the value of new information varies with ξ.

If the seller's learning affects prices through only 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 additional mechanisms due to 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.

$(E_{y^{t+1}|y^t} V(\mu(y^{t+1}), \sigma_{t+1}, K, p) - V(\mu(y^t), \sigma_t, K, p))$, 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)).

[23] 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 is more likely to adjust her belief about ξ downward because of the dynamic adverse-selection effect, it is possible for her 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 alternative 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). Overall, these predicted pricing patterns are consistent with what we observe in the data (see Section 3.2).

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 intertemporal pattern in the sale probabilities to identify the variance of ξ_j.[24]

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.[25] Suppose that we have the following reduced-form pricing equation for

[24] As an alternative estimation approach to deal with the endogeneity problem, one may estimate the demand-side model using the pseudo-policy function approach (e.g., Erdem et al. (1999); Ching (2010a); Ching and Ishihara (2010); Ishihara and Ching (2016)).

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

the initial price:

$$p_{j1} = Z_j \varphi + \zeta_{j1}. \tag{2}$$

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:

$$v_{j1} = X_j\beta + \alpha p_{j1} + \psi\zeta_{j1} + \eta_j + \varepsilon_{j1}, \tag{3}$$

where $\psi = \frac{\sigma_{\xi\zeta}^2}{\sigma_\zeta^2}$. It is worth noting that we use the linear pricing equation as an approximation of the optimal pricing policy. The simplification allows us to use the control function method to separately estimate the demand model.

Without loss of generality, let $\varepsilon_{0t} = 0$ because all that matters for the likelihood of sales data is the distribution of $\varepsilon_{jt} - \varepsilon_{0t}$. For identification, we normalize the variance of ε_{jt} to be one. The parameters we need to estimate for the demand model are, therefore, $\theta_d \equiv (\beta, \alpha, \psi, \sigma_\eta)$.

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 $\theta_{d1} \equiv (\beta, \alpha, \psi)$—i.e., all the demand-model parameters except for σ_η. 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, which is similar in spirit to the instruments used in Berry et al. (1995).[26] The inventory level is a valid excluded variable because the 'within-dealer' competition across similar cars should affect the prices, p_{j1}, but it should not directly affect buyer 1's utility for car j.

Then, we estimate the demand as a probit model. The sale probability of car j on the first day,

[26]The instruments used in Berry et al. (1995) include a constant, the number of own-firm products and the number of rival-firm products. The number of own-firm products excludes the car model under consideration. They discuss these instruments in equation (5.8) on page 861: "Note also that one of our characteristics is a constant term, so that the number of own-firm products and rival-firm products become instruments."

h_{j1}, can be written as follows:

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

where $P(\eta_j + \varepsilon_{j1})$ is the probability measure of the normal distribution with the mean as zero and the covariance as $1 + \sigma_\eta^2$ (note that $\eta_j \perp \varepsilon_{j1}$ by assumption).

We focus on the first day in the first step of our estimation because, the joint distribution of ξ_j and ζ_{jt} changes as the price changes for some cars over time. We 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:

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

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 on each car to estimate σ_ξ. Let $\tau_j \in \{1, 2, 3, ...\}$ be the number of days that car j takes 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:

$$\tilde{h}_{jt} \equiv \Pr\left(I_{jt} = 1|X_j, p_{jt}, \hat{\zeta}_{jt}, \eta_j; \hat{\theta}_{d1}\right),$$
$$= \int 1\left\{v_{jt} > 0\right\} d\Phi\left(\varepsilon_{jt}\right),$$

where $\Phi\left(\varepsilon_{jt}\right)$ is the probability measure of the standard normal distribution. Thus, we have the following expression for the likelihood of the observation of car j:

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

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:

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

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 subsection.

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.[27] 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.[28]

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

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

[28] 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.

depend on θ_s, yields the partial likelihood of the observed price sequence of car j, $\tilde{l}\left((p_{jt}, I_{jt})_{t=1}^{T_j} | \theta_s\right)$:

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

The expression after the second equality above makes it explicit that 1) 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 2) we, as researchers, do not observe the information that 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 θ_s using the MPLE method as follows:

$$\begin{aligned}\hat{\theta}_s = \arg\max_{\theta_s}\sum_{j=1}^{J}&\left[\log l\left(p_{j1}|\theta_s\right) + \right.\\ \sum_{t=2}^{T_j}\log\int l\left(p_{jt}|y_j^t, p_{j,t-1}; \theta_s\right) f&\left.\left(y_j^t|\left(p_{j\tau}, I_{j\tau}\right)_{\tau=1}^{t-1}; \theta_s\right) dy_j^t\right].\end{aligned} \tag{4}$$

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—parameterizing the value function using the Chebyshev polynomials and iterating over the policy function until convergence. Appendix B contains the details of the numerical solution method.

Computing the log-likelihood function in (4) is difficult because it involves high-dimensional integrations over the conditional distributions of serially correlated signals.[29] 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.[30] The details

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

[30]See Rubin (1988). See, also, Fernandez-Villaverde and Rubio-Ramirez (2007), Flury and Shephard (2011) and Gallant et al. (2018) for examples of applying SIR to simulate the likelihood function when estimating dynamic

of computing the likelihood are in Appendix C.

Identification

We do not estimate the daily discount factor δ, but calibrate it to match an annual rate of 25%.[31] 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 higher initial price with smaller variance (after controlling for the observable car attributes). 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.[32]

The distribution of one-time price adjustments on any given day (e.g., price changes on the second day relative to the first day) helps jointly identify σ_s and ϕ_1. To see this, first assume that the realized menu cost on a given day t is zero. Then, the probability of no price changes would be zero. Given the seller's prior belief about ξ, characterized by $(\sigma_{t-1}, \mu_{t-1})$, σ_s determines the distribution of μ_t, and, consequently, the distribution of price adjustments on day t. Specifically, because $\mu_t - \mu_{t-1} = \sigma_{t-1}^2/(\sigma_{t-1}^2 + \sigma_s^2) * (y_{t-1} - \mu_{t-1})$, the magnitude of $\mu_t - \mu_{t-1}$ increases as σ_s decreases, for any given signal y_{t-1}. As a result, with a smaller σ_s, the magnitude of the price adjustment, $p_t - p_{t-1}$, is also larger. When the realized menu cost is positive, however, price is adjusted only when the resulting positive impact on the expected payoff is greater than the menu cost. Because smaller price adjustments generally have smaller impacts on the expected payoff, any realized level of the menu cost tends to censor relatively small price adjustments but not larger ones. To demonstrate this, Figure 4 plots the predicted distributions of price adjustments (using

models. Alternatively, one can express $\tilde{l}\left((p_{jt}, I_{jt})_{t=1}^{T_j} | \theta_s\right)$ as follows:

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

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

[31] 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.

[32] It is worth noting that it is not possible to separately identify σ_ξ and σ_0 from the supply-side information alone. As explained in the previous section, σ_ξ is identified and estimated from the demand side first.

the estimated model) on the second day for different levels of menu costs. It shows that the menu cost has a large impact on the section around zero but little impact on the two end-sections of the distribution. Therefore, roughly speaking, the two end-sections of the distribution of one-time price adjustments on any given day are the main source of identification for σ_s. With σ_s identified, the probability of no price adjustments helps identify the menu cost, ϕ_1.

Lastly, the effect of inventory on price, captured by parameter ϕ_2, is identified by the residual impact of K_{jt} on p_{jt} that cannot be rationalized by the predicted responses of prices to the competition effect.

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 2005-2009, dropping older cars. We focus on a small set of models with similar observable characteristics, so it is reasonable to assume that ξ is independent of car model and other observable attributes. Cars with missing data for any attributes or price are also dropped. For each car, we use the first day's data in the first step of the estimation, and 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.

There are 975 unique cars in the demand-estimation sample: 12.5% are sold on the first day, and 9.6% of the remaining cars are sold on the second day. Table 8 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 Accords, the model with the largest number of cars in our data, that appeared in 2011 in the above sample. Table 9 presents the summary statistics of this subsample of 154 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.[33] Because CarMax's observed pricing and sales patterns do not vary substantially across car models (c.f. Table 6), we are likely to get similar estimates with 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. Then, 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. Lastly, we discuss the managerial implications of our results.

6.1 Model Estimates

Demand Model

Table 10 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 coefficient of the current inventories is -0.055, statistically significant at the 1% level. The estimate is consistent with the seller's incentive to avoid stock-out and with the fact 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, those with more powerful engines and those with lower mileage. It is worth noting that the adjusted R-squared of the regression is 0.74, suggesting that there is still a fair amount of price variation that the variations in the observable attributes cannot explain.

Column (2) of Table 11 presents the parameters of the demand model, estimated using the adapted control function approach. The estimated price coefficient is -0.793 and is statistically

[33] It takes around eight days to estimate the dynamic pricing model using the subsample of Honda Accords. We obtained similar, but statistically less significant, parameter estimates when using other models with smaller numbers of observations.

significant at the 5% level. The estimated coefficients of other car attributes show that buyers' values for the cars increase with the engine volume (significant at the 5% level) and decrease with mileage (significant at the 5% level). The coefficient of the price residual, ζ, is 0.593 (with the p-value being 12%), indicating a positive correlation between ξ and ζ.

The estimated standard deviation of η is 0.439, 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 0.90. In comparison, the standard deviation of $X_j\beta$, the utility of observable attributes, is estimated to be 1.62. Thus, the across-car variation in latent car quality is comparable in importance to that of the observable attributes.

Column (1) of Table 11 presents the estimated parameters of the demand model without using the control function approach or the instrumental variable. A comparison with Column (2) shows that the estimated price elasticity using the instrument is significantly larger, demonstrating the importance of instrumenting for price.

Dynamic Pricing Model

Table 12 presents the estimates of the structural parameters of the dynamic pricing model. Because it is computationally very challenging to account for the variances of the estimated parameters in the demand model, we treat the estimated demand model as the true demand model when computing the standard errors for parameters in the pricing model (see, for example, Benkard (2004), Ching (2010b) and Ching (2010a) for papers taking a similar approach). However, we do not expect the adjustment to materially affect our results because most of the parameters in the pricing model are estimated with great precision.

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 at 0.46 and 0.69, respectively. Recall that the demand estimation shows that the seller's ex ante belief about ξ has a standard deviation of 0.90. These results suggest that: (1) CarMax's initial assessment leads to information about car-specific demand that is much more accurate; and (2) the learning in the subsequent selling process is also quite informative.

The estimated mean of the menu cost, ϕ_1, is \$181, 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 -0.10 (though not statistically significant). This is consistent with the seller's objective of avoiding stock-outs, as explained earlier. The holding cost is estimated at \$41 per day, which seems reasonable given the maintenance cost and the opportunity cost of the parking space (when filled 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 the 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 on all the cars in the estimation sample.

The top panel of Table 13 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, \$17,861 and \$17,577, and those predicted by the model are, respectively, \$17,657 and \$17,437. The predicted average prices are close to those in the data, with the respective relative prediction errors of 1% and 0.8%. Second, in the estimation sample, 50.0%, 80.5% and 92.2% of the cars are sold within, for example, the first six, 15 and 30 days, respectively. With the prices predicted by the pricing model, the demand model predicts that 51.9%, 80.8% and 95.1% of the cars are sold within the corresponding time frames. Thus, the predicted sale probabilities are also very close to the actual data.

The above 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 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 intertemporal price movement. The bottom panel of Table 13 shows that, in the estimation sample, 26.6% of the cars experience total-price drops and 0.6% experience total-price increases,[34] while the corresponding model predictions are 22.4% and 8.6%. The average magnitudes of the total-price drops and increases are, respectively, $524 and $300 in the estimation sample, while the corresponding model predictions are $558 and $426.

Thus, our model is able to produce the overall downward trend in price adjustments, predicting many more price decreases than price increases. The model-predicted likelihood of and average magnitude of total price drops are close to those in the data. The model, however, predicts a significantly greater likelihood of price increases and a larger magnitude of increases than the data show. The significant overprediction of the total-price increases could be a limitation of the Gaussian learning process that we assume in the pricing model. In particular, the signals in the Gaussian learning model have support on the entire real line. The large support of signals may make it more likely, than in reality, for the seller to receive signals positive enough to trigger price increases. This seems a price necessary for the tractability for our 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 in 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 show 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

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

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. With no initial assessment, the seller's belief about ξ when setting the initial price is just $N(0, \sigma_\xi^2)$. In the scenarios involving "weak learning," we set the standard deviation of the signals received in the selling process, σ_s, at three times the original estimate.[35] To provide a benchmark, we also simulate the counterfactual scenario in which the seller has perfect information about ξ from the beginning.

Table 14 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 received (at least to some extent) the benefit of subsequent learning in the selling process. Given our discussion 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 collected information can lead to a) a higher expected transaction price; b) a lower total holding cost due to a shorter time until sale; and c) a smaller total menu cost due to less need to adjust

[35] 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. In particular, in 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 removed from 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.

prices. Our simulation results show that, relative to "Learning without Assessment," adding the initial assessment increases the average expected transaction price by $92, lowers the total holding cost by $20, and reduces the total menu cost by $25.5. In total, the estimated value of the initial assessment is approximately $139 per car.

The seller may benefit from subsequent learning in the selling process because it can lead to faster adjustments of the list 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 lowers the total holding cost by $241. The benefit is achieved with an additional menu cost of $11.5 and a decrease in the transaction price of $26. Overall, the estimated value of the subsequent learning is about $203 per car. Table 15 summarizes the value of both the initial assessment and subsequent learning.

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.[36] Therefore, these results show that subsequent learning in the selling process can greatly improve the profitability of dealers who conduct no initial assessment for their cars; however, even with subsequent learning, a careful inspection of cars and research into local consumers' preferences can still benefit the seller significantly.

Note that, in general, the initial assessment and the subsequent learning are substitutes: the marginal value of one is lower in the presence of the other. The results reported in Table 15 show that, when there is only weak learning, adding the initial assessment increases the net revenue by $297, 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 $45, which is less than the value of the subsequent learning estimated above.

Table 14 also shows that, relative to "Assessment and Learning," the average expected net revenue can be further increased by $151 if the seller begins with perfect information about car-specific demand. That improvement, compared to what has already been achieved overall, seems

[36] 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 2011 annual report. The reported average selling price of used cars was $18,019 (page 22). CarMax sold 396,181 used vehicles in 2011 (page 22). The total earnings before taxes for used vehicles was 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 was, therefore, 6.9%, calculated as the ratio of the net earnings before tax per car to that of the average selling price.

relatively modest.

Finally, we investigate the extent to which the menu cost prevents the seller from effectively utilizing the information learned in the selling process. Specifically, we simulate the counterfactual scenario of "Half Menu Cost," which modifies "Assessment and Learning" by lowering the mean of the menu cost by half. Table 16 shows that, with the menu cost cut by half, the average expected transaction price decreases by \$10; and the changes in the average total menu cost incurred per car and time to sell are also small. As a result, the average expected net revenue increases by only \$4 per car.

The above results suggest that the menu cost (with the mean of \$181) seems to have a limited impact on 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 that the seller actually incurs is small: in the case of "Assessment and Learning," for example, with 27.2% of the cars experiencing at least one price adjustment, the average total menu cost per car is only \$5.2. Therefore, as long as the seller is not too impatient, the menu cost will not significantly impact her performance.

6.4 Implications for Managers in the Used-Car Retail Market

Although most dealers conduct some inspection and make price adjustments, there seems substantial room for them to improve pricing with more-accurate initial assessment and effective subsequent learning. Anecdotal evidence—from our conversations with a co-founder of CarMax and pricing practitioners in the market—suggest that CarMax holds a significant advantage over other dealers due to its ability to use information technology in setting and adjusting prices. Our empirical findings in across-dealer comparisons also suggest less-sophisticated pricing by other dealers. In particular, Table 5 shows that, relative to CarMax, other dealers in our data needed significantly larger total price adjustments and more time to sell their cars. We also find that other dealers were slower to adjust their prices (which is not reported in the paper to save space). These empirical findings suggest that other dealers were 1) basing their starting prices on less-accurate initial assessments; and 2) not as effective in utilizing new information to guide dynamic price adjustments.

When dealers consider whether to switch to a more information-based pricing approach, the benefits of such a pricing strategy must be weighed against the costs that dealers have to incur. Though the variable cost of thoroughly assessing each car should be lower than the estimated value of the initial assessment, building up the necessary inspection facilities, databases and IT infrastructure involves large investments and can take multiple years. Given the substantial fixed costs, larger dealers are likely to be in a better position to exploit an information-based pricing strategy due to their greater economies of scale.

In light of the associated costs and our estimates of the relative payoffs to improving initial assessments ($139) vs. subsequent learning ($203), it seems advisable for dealers to take a multiple-step strategy to transition to (or experiment with) a more information-based pricing strategy. In particular, it can be most cost-effective for dealers to start by improving price-adjustment efficiency through more effective uses of new information readily accessible in the selling process. In the beginning, dealers can work on developing methods to systematically record information revealed by no-sale events and their interactions with buyers. They can then build basic statistical models, incorporating this information into their demand forecasts, and making their price adjustment decisions more tightly based on their latest demand forecasts (Footnote 2 includes some details on how, technically, CarMax was able to collect information at various points of the selling process and incorporate such information into its pricing decisions).

After this initial step, with the building-up of inspection facilities, databases of historical transactions and other market intelligence, dealers should be more ready to enhance their initial assessments by putting their cars through more thorough inspections. They then would be able to incorporate those inspection results into their demand forecast models and fine-tune their demand forecasts as they grow their database on historical transactions and local consumers. Dealers could also work on further streamlining the information collection and coding process and refining their pricing algorithms to accommodate their more accurate demand forecasts.

Managers should expect some changes in price trajectories after a transition to information-based pricing. A more accurate initial assessment would lead to better-customized initial prices for individual cars. In general, the initial prices would be impacted in two ways. First, there would be larger variance in initial prices. This is a direct implication of more information being incorporated into initial prices. Second, initial prices should be set lower, on average. This is because, with more accurate initial assessment and less uncertainty left to be resolved later, the option value of selling a car in the future becomes lower. Therefore, with the switch to an information-based pricing strategy, managers should expect these two effects on initial prices.

Though our policy suggestions speak directly to dealers in the used-car retail market, they may also be of interest to sellers of other products with substantial item-specific heterogeneity.

7 Summary and Concluding Remarks

The main challenge for dealers in pricing products with substantial item-specific heterogeneity comes from their ex ante lack of information about the demand for individual items. In this paper, we develop a structural model of dynamic pricing for this type of product, 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 prices for individual cars.

We estimate 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 these products.

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

We treat 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 also be informative about the demand for another item. We leave the study of this phenomenon to future research.

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

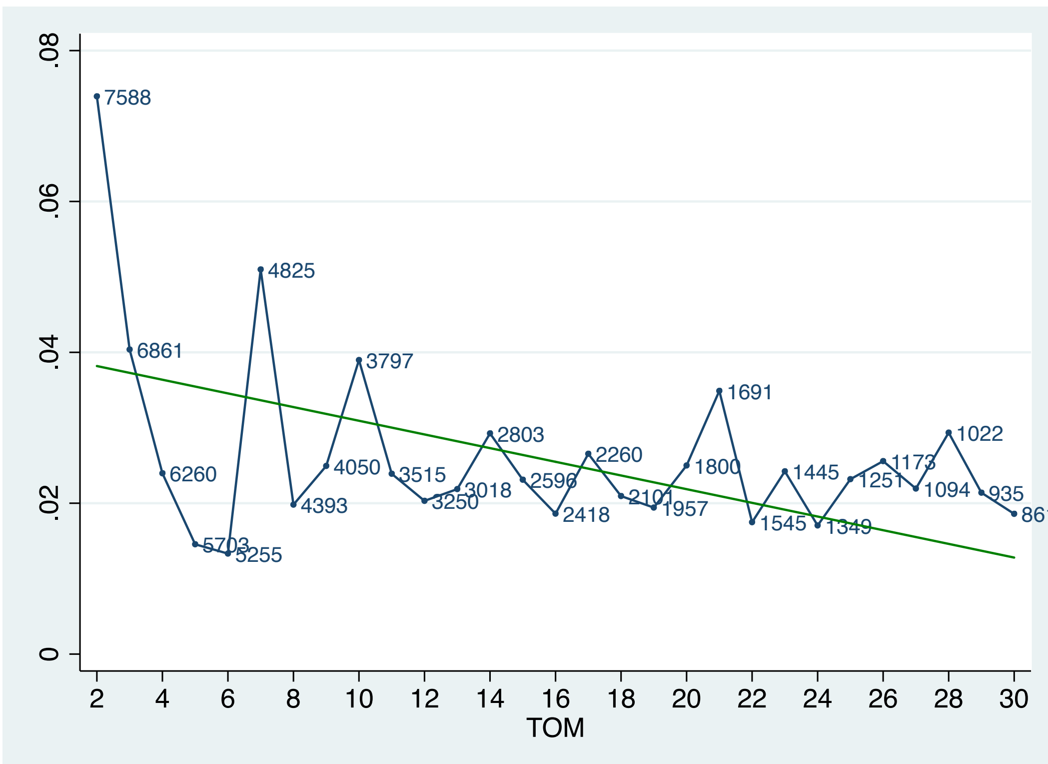

Note: For each day of time on the market (TOM), the dot represents the percentage of remaining cars with adjusted prices relative to the previous day, and the accompanying number is the number of cars that remain on the market on that particular day.

Figure 3: Frequency of Price Adjustments by Time on the Market: Other Dealers

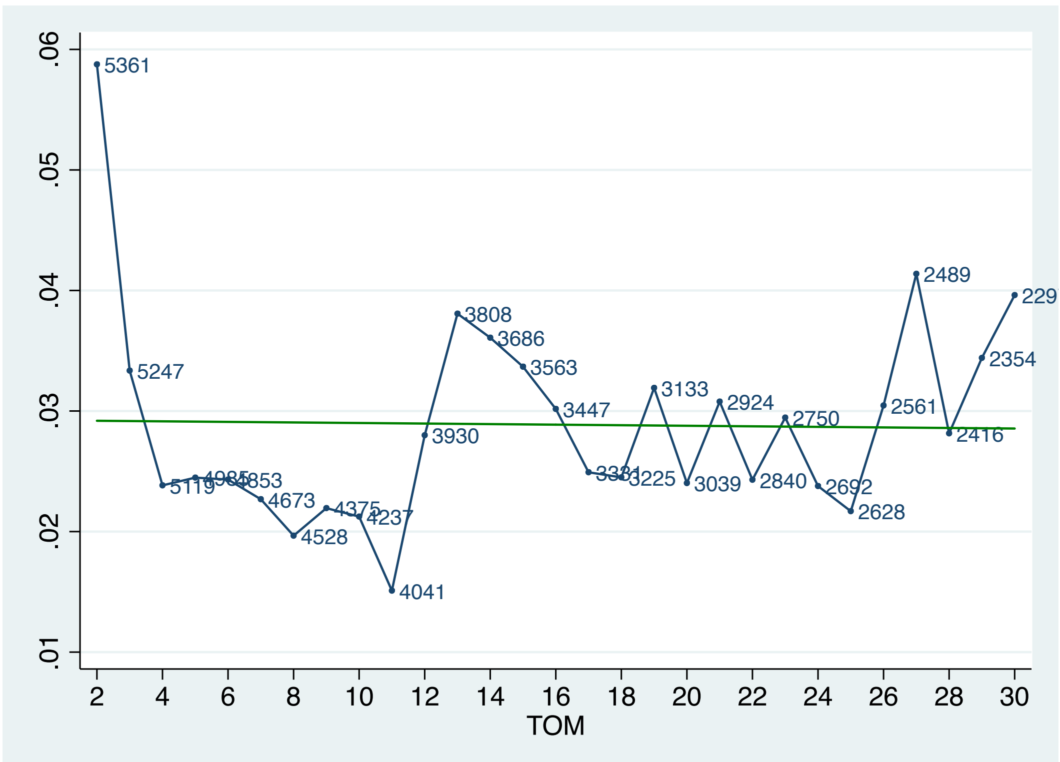

Note: This figure replicates Figure 2 using data from the five other largest local dealers from the same period. For each day of time on the market (TOM), the dot represents the percentage of remaining cars that have their prices adjusted relative to the previous day, and the accompanying number is the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

Figure 4: The Impact of Menu Cost on the Distribution of One-time Price Adjustments

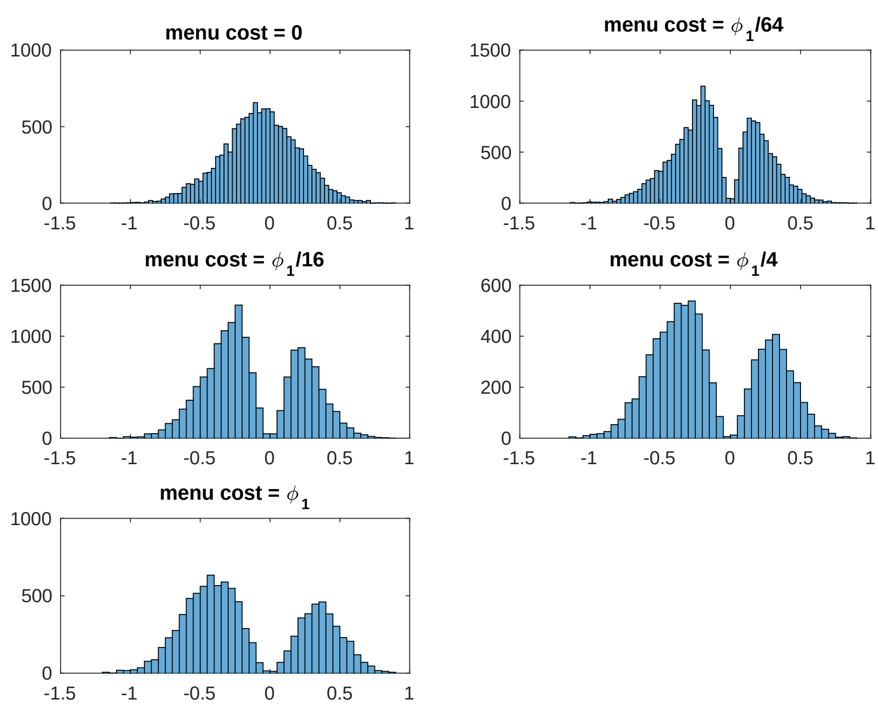

Note: The distribution is simulated using the estimated dynamic pricing model.

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: The Evolution of Inventory K_t and the Focal Car's Price

	Dependent Variable: K_t			
	OLS	OLS	Car Fixed Effects	Car Fixed Effects
K_{t-1}	0.950***	0.950***	0.759***	0.759***
	(0.00427)	(0.00431)	(0.0138)	(0.0141)
p_{t-1}		0.00503		0.155
		(0.00444)		(0.126)
Constant	0.255***	0.172*	1.271***	-1.332
	(0.0199)	(0.0741)	(0.0735)	(2.100)
N	8406	8406	8406	8406
adj. R^2	0.907	0.907	0.583	0.583

Standard errors (clustered at the car level) in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 8: 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 2010 and 2011. These models are Honda Accord and Civic, Nissan Altima, Toyota Camry and Corolla, and Hyundai Sonata.

Table 9: Estimation Sample (Dynamic Pricing Model) Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Price ($1,000)	17.86	2.25	12.7	22	154
Mile (10k)	3.66	2.08	0.62	9.99	154
Engine volume (liter)	2.57	0.38	2.4	3.5	154
Wheelbase	109.14	0.84	107	110	154
Model year 06	0.06	0.25	0	1	154
Model year 07	0.19	0.4	0	1	154
Model year 08	0.45	0.5	0	1	154
Model year 09	0.24	0.43	0	1	154
Seller's inventory of the same model	3.77	1.87	1	9	154
Total price increases ($1,000)	0.30	0	0.30	0.30	1
Total price decreases ($1,000)	0.52	0.31	0.10	1.40	41

Note: The estimation sample of the dynamic pricing model includes all the Honda Accord cars that the CarMax store in our data carried in 2011.

Table 10: Reduced-form Pricing Equation

Variable	Coefficients	Std. Err.
Constant	-9.757	7.931
Mile (10k miles)	-0.544**	0.027
Engine volume (liter)	2.30***	0.128
Wheelbase	0.195***	0.073
Inventory	-0.055***	0.015
Dummies for car model, model year, month and year	Yes	
Number of observations	975	
Adj. R^2	0.735	

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 11: Probit Demand Model Parameter Estimates

Variables	Without IV	With IV
Constant	-2.207	-0.535
	(12.329)	(12.954)
Price ($1,000)	-0.203***	-0.793**
	(0.044)	(0.382)
Mile (10k miles)	-0.112***	-0.429**
	(0.046)	(0.204)
Engine volume (liter)	0.386*	1.726**
	(0.205)	(0.882)
Wheelbase	0.032	0.077
	(0.114)	(0.135)
Price residual (ζ)		0.593
		(0.382)
Std. dev. of η:	0.410***	0.439***
	(0.077)	(0.007)
Dummies for car model, model year, year, month	Yes	
Number of unique cars	975	

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

Table 12: Dynamic Pricing Model Parameter Estimates

Parameters	Coefficients	Std. Err.
Std. dev. of the initial assessment signal (σ_0)	0.46***	0.01
Std. dev. of signals received after the initial assessment (σ_s)	0.69***	0.03
Mean of the menu cost (ϕ_1)	$181.37***	21.85
The inventory effect (ϕ_2)	-0.10	0.23
Holding cost (ϕ_3)	$40.69***	1.48
Num. of cars	154	

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

Table 13: Model Fit

Performance metrics	Data	Model prediction
Price and sales levels		
Average initial price	17,861	17,657
Average transaction price	17,577	17,437
Cars sold in the first six days (%)	50.0	51.9
Cars sold in the first 15 days (%)	80.5	80.8
Cars sold in the first 30 days (%)	92.2	95.1
Inter-temporal price adjustments		
Cars with total price decrease (%)	26.6	22.4
Average total price decrease	524	558
Cars with total price increase (%)	0.6	8.6
Average total price increase	300	426

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 equivalent to 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 14: Policy Experiment Results

Scenarios	Net revenue	Transaction price	Time to Sell	Total menu cost
Assessment and Learning	17,082	17,437	9.6	5.2
Learning w/o Assessment	16,943	17,345	10.1	30.7
Assessment and Weak Learning	17,037	17,512	12.6	1.3
Weak Learning w/o Assessment	16,740	17,371	16.0	19.2
Perfect Initial Information	17,233	17,609	10.3	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 equivalent to 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 15: The Values of the Initial Assessment and Subsequent Learning

(Assessment, Learning)	*Yes*	*Weak*	Value of subsequent learning
Yes	17,082	17,037	
No	16,943	16,740	203
Value of initial assessment	139		

Note: The values reported for the four scenarios are the average expected net revenues taken from Table 14.

Table 16: The Impact of Menu Cost

Scenarios	Net revenue	Transaction price	Time to Sell	Total menu cost
Assessment and Learning	17.082	17,437	9.6	5.2
Half Menu Cost	17,086	17,427	9.3	4.7

Note: The scenario "Assessment and Learning" is the same as that in Table 14. 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. *Working Paper*.

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. *Working Paper*.

Benkard, L. (2004). A dynamic analysis of the market for wide-bodied commercial aircraft. *Review of Economic Studies 71*, 581–611.

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. T. (2010a). Consumer learning and heterogeneity: Dynamics of demand for prescription drugs after patent expiration. *International Journal of Industrial Organization 28*(6), 619–638.

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

Ching, A. T., T. Erdem, and M. P. Keane (2013). Invited paper-learning models: An assessment of progress, challenges, and new developments. *Marketing Science 32*(6), 913–938.

Ching, A. T., T. Erdem, and M. P. Keane (2017). Empirical models of learning dynamics: A survey of recent developments. In *Handbook of Marketing Decision Models*, pp. 223–257. Springer.

Ching, A. T. and M. Ishihara (2010). The effects of detailing on prescribing decisions under quality uncertainty. *Quantitative Marketing and Economics 8*(2), 123–165.

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.

Einav, L., C. Farronato, J. Levin, and N. Sundaresan (2018). Auctions versus posted prices in online markets. *Journal of Political Economy 126*(1), 178–215.

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.

Erdem, T., M. P. Keane, and B. Sun (1999). Missing price and coupon availability data in scanner panels: Correcting for the self-selection bias in choice model parameters. *Journal of Econometrics 89*, 177–196.

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 (2011). Bayesian inference based only on simulated likelihood: particle filter analysis of dynamic economic models. *Econometric Theory 27*(5), 933–956.

Gallant, A. R., H. Hong, and A. Khwaja (2018). A bayesian approach to estimation of dynamic models with small and large number of heterogeneous players and latent serially correlated states. *Journal of Econometrics 203*(1), 19–32.

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.

Handel, B. R. and K. Misra (2015). Robust new product pricing. *Marketing Science 34*(6), 864–881.

Hitsch, G. J. (2006). An empirical model of optimal dynamic product launch and exit under demand uncertainty. *Marketing Science 25*(1), 25–50.

Ishihara, M. and A. T. Ching (2016). Dynamic demand for new and used durable goods without physical depreciation: The case of japanese video games. *Working Paper*.

Larsen, B. (2014). The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. *Working Paper*.

Lewis, G. (2011). Asymmetric information, adverse selection and online disclosure: The case of ebay motors. *American Economic Review 101*(4), 1535–46.

Mason, R. and J. Välimäki (2011). Learning about the arrival of sales. *Journal of Economic Theory 146*(4), 1699–1711.

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-Magné, and J. Rust (2015). The home selling problem: Theory and evidence. *International Economic Review 56*(2), 457–484.

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.

Newberry, P. W. (2016). An empirical study of observational learning. *The RAND Journal of Economics 47*(2), 394–432.

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.

Tadelis, S. and F. Zettelmeyer (2015). Information disclosure as a matching mechanism: Theory and evidence from a field experiment. *American Economic Review 105*(2), 886–905.

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. *Working Paper*.

Zhang, J. (2010). The sound of silence: Observational learning in the us kidney market. *Marketing Science 29*(2), 315–335.

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:

$$\begin{aligned}
&\max_{(p_t)_{t=1}^{\infty}} \left\{ E_{\xi} E_{(y^t)_{t=1}^{\infty}|\xi} \left(E_{(K_t)_{t=2}^{\infty}|K_1} \sum_{t=1}^{\infty} \delta^{t-1} \chi_t E_{\varphi_t} \pi_t \left(p_t, \xi, K_t, \varphi_t\right) \right) \right\} \\
&= E_{y^1} \max_{(p_t)_{t=1}^{\infty}} \left\{ E_{\xi|y^1} E_{(y^t)_{t=2}^{\infty}|\xi, y^1} \left(E_{(K_t)_{t=2}^{\infty}|K_1} \sum_{t=1}^{\infty} \delta^{t-1} \chi_t E_{\varphi_t} \pi_t \left(p_t, \xi, K_t, \varphi_t\right) \right) \right\},
\end{aligned} \tag{5}$$

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:

$$\begin{aligned}
&\max_{(p_\tau)_{\tau=t}^{\infty}} \left\{ E_{\xi|y^t} E_{(y^\tau)_{\tau=t+1}^{\infty}|\xi, y^t} \left(E_{(K_\tau)_{\tau=t+1}^{\infty}|K_t} \sum_{\tau=t}^{\infty} \delta^{\tau-t} \chi_\tau E_{\varphi_\tau} \pi_\tau \left(p_\tau, \xi, K_\tau, \varphi_\tau\right) \right) \right\} \\
&= \max_{p_t} \Bigg\{ E_{\xi|y^t} E_{\varphi_t} \pi_t \left(p_t, \xi, K_t, \varphi_t\right) + \max_{(p_\tau)_{\tau=t+1}^{\infty}} E_{\xi|y^t} \{ \left(1 - D\left(p_t, \xi\right)\right) \\
&\quad \delta E_{(y^\tau)_{\tau=t+1}^{\infty}|\xi, y^t} E_{(K_\tau)_{\tau=t+1}^{\infty}|K_t} \sum_{\tau=t+1}^{\infty} \delta^{\tau-(t+1)} \chi_\tau E_{\varphi_\tau} \pi_\tau \left(p_\tau, \xi, K_\tau, \varphi_\tau\right) \} \Bigg\} \\
&= \max_{p_t} \{ E_{\xi|y^t} E_{\varphi_t} \pi_t \left(p_t, \xi, K_t, \varphi_t\right) + E_{\xi|y^t} \left(1 - D\left(p_t, \xi\right)\right) E_{y^{t+1}|\xi, y^t} E_{K_{t+1}|K_t} \\
&\quad \delta \max_{(p_\tau)_{\tau=t+1}^{\infty}} E_{\xi|y^{t+1}} \left\{ E_{(y^\tau)_{\tau=t+2}^{\infty}|\xi, y^{t+1}} E_{(K_\tau)_{\tau=t+2}^{\infty}|K_{t+1}} \sum_{\tau=t+1}^{\infty} \delta^{\tau-(t+1)} \chi_\tau E_{\varphi_\tau} \pi_\tau \left(p_\tau, \xi, K_\tau, \varphi_\tau\right) \right\} \},
\end{aligned} \tag{6}$$

where the second equality follows by the Law of Iterated Expectations. Taken together, reformulations (5) and (6) 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:

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

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

where $S_t \equiv \left(\left(\mu\left(y^t\right), \sigma_t\right), p_{t-1}, K_t\right)$. Among the state variables, $\left(\mu\left(y^t\right), \sigma_t, p_{t-1}\right)$ 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 $D(p) \equiv D(p, \xi)$, 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:

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

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:

$$V^*(S_t) = F(\bar{\varphi}(S_t|V^*))\, W(S_t|V^*) + (1 - F(\bar{\varphi}(S_t|V^*)))\, U(S_t|V^*) - \int_0^{\bar{\varphi}(S_t|V^*)} \varphi_t dF(\varphi_t). \qquad (7)$$

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 $S_{t+1}(p) \equiv \left(\left(\mu\left(y^{t+1}\right), \sigma_{t+1}\right), p, K_{t+1}\right)$.[37] 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:

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

[37] This equation is an analog of equation (7), with $(V^*, \bar{p}(S_t|V^*), \bar{\varphi}(S_t|V^*))$ replaced by $(V_k, \bar{p}_{k-1}(S_t), \bar{\varphi}_{k-1})$

where:

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

If we pick a grid of N different points of S_t: $(S_1, ..., S_N)$, then we can solve for Ψ by using the least squares criterion as follows:

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

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

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

To implement the algorithm, we start with an initial guess of $(\bar{p}_0(S_t), \bar{\varphi}_0(S_t))$, 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 (4). 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 $\xi_j \sim N\left(0, \sigma_\xi^2\right)$, $\epsilon_{j0} \sim N\left(\xi_j, \sigma_0^2\right)$ and $\epsilon_{jt} \sim N\left(\xi_j, \sigma_s^2\right)$ 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 $f\left(\xi_j|y_j^t\right) = \frac{1}{\sigma_t}\phi\left(\frac{\xi_j - \mu_{jt}}{\sigma_t}\right)$, 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 $f\left(y_{jt}|y_j^t\right) = f(y_{jt}|\mu_{jt}) = \frac{1}{\sqrt{\sigma_t^2 + \sigma_s}}\phi\left(\frac{y_{jt} - \mu_{jt}}{\sqrt{\sigma_t^2 + \sigma_s}}\right)$ 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:

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

where $\mu_t(p_{jt})$ is defined by $p_{jt} = \wp_0(\mu_t, \sigma_t, K_{jt})$ and $f_{\mu_{j1}}(\mu_1) = \frac{1}{\frac{\sigma_\xi^2}{\sqrt{\sigma_\xi^2 + \sigma_0^2}}}\phi\left(\frac{\mu_1}{\frac{\sigma_\xi^2}{\sqrt{\sigma_\xi^2 + \sigma_0^2}}}\right)$. 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.[38] 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

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

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:

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

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

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

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:

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

[38] 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 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 $\left(\mu_{jt}\right)_{s=1}^{ns}$. It is sufficient for us to keep the record of the simulated random sample in the form of $\left(y_{jt,s}, \mu_{jt,s}\right)$ 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: resample 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.