

Inventories and the automobile market

Adam Copeland*

Wendy Dunn**

and

George Hall***

This article studies the within-model-year pricing, production, and inventory management of new automobiles. Using new monthly data on U.S. transaction prices, we document that, for the typical vehicle, prices fall over the model year at a 9.0% annual rate. Concurrently, both sales and inventories are hump shaped. To explain these time series, we formulate an industry model for new automobiles in which inventory and pricing decisions are made simultaneously. The model predicts that automakers' build-to-stock inventory management policy substantially influences the time series of prices and sales, accounting for four tenths of the price decline observed over the model year.

1. Introduction

■ Two common features of durable goods markets are high levels of inventories relative to sales and declining prices over the product cycle. There has been substantial research explaining why durable goods prices decline, with most existing theories focusing on intertemporal price discrimination (e.g., Stokey, 1979; Conlisk, Gerstner, and Sobel, 1984) or fashion (e.g., Lazear, 1986; Pashigian, 1988; Pesendorfer, 1995). We build on this body of work by emphasizing the significant role that firm-held inventories can play in explaining price declines.

We focus on automakers, the quintessential durable goods producer. We begin by constructing a monthly data set on transaction prices, sales, and inventories for Big Three vehicles. We document four stylized facts: (i) average retail prices, net of rebates and incentives, decline by

*Federal Reserve Bank of New York; adam.copeland@gmail.com.

**Board of Governors of the Federal Reserve System; wendy.e.dunn@frb.gov.

***Brandeis University; ghall@brandeis.edu.

We thank Ana Aizcorbe, Steve Berry, Andrew Cohen, Gautam Gowrisankaran, Amil Petrin, John Rust, John Stevens, and participants at numerous conferences and seminars for their helpful comments. We also received valuable comments from the editor and three referees. Finally, we thank Bob Schnorbuss for helping us obtain and interpret the data from J.D. Power, and Associates. Hall gratefully acknowledges financial support from the Alfred P. Sloan Foundation. The views expressed in this article are those of the authors and not necessarily reflective of views at the Board of Governors, the Federal Reserve Bank of New York, or the Federal Reserve System.

9% at an annual rate; (ii) for about half the calendar year, automakers simultaneously sell two vintages of the same model, during which the older vintage sells for a 9% discount; (iii) sales and inventories are humped shaped over the product cycle; and (iv) the mean ratio of inventories to sales is 75 days.

To explain these stylized facts, we develop and parameterize a two-sided industry model. We describe the firm as a dynamic inventory control problem. The joint production/pricing decision we model is a classic issue in the operations research literature going back to Whiten (1955) and Karlin and Carr (1962).¹ We extend the theory by allowing the firm to sell two vintages simultaneously, which is a frequent occurrence in the automobile industry (the second stylized fact). Further, because Hamermesh (1989) and Bresnahan and Ramey (1994) document that durable goods manufacturers frequently adjust their rate of production by shutting down the plant for a week or two at a time, we incorporate the nonconvex cost structure from Hall (2000) into the model to induce our fictional producer to mimic the observed production-bunching behavior.

On the consumer's side, we estimate preferences for automobiles by employing the econometric methodology developed in the discrete-choice literature (for example, Berry, Levinsohn, and Pakes, 1995; Goldberg, 1995; Petrin, 2002; to name a few). Our approach differs from the standard one in three ways. First, motivated by Kahn (1987, 1992), who finds that inventories are productive in generating greater sales at a given price, we include an inventory-based measure of variety in the consumer's indirect utility function. This allows us to compute by how much demand changes when manufacturers, by altering their stock of inventories, change the variety they offer consumers. Second, we estimate our demand-side model at a quarterly, rather than an annual, frequency using transaction rather than list prices; thus, we can estimate how the demand curve shifts throughout the model year. Third, our data let us differentiate multiple vintages of the same model. Hence, we allow consumers to choose among multiple vintages within and across models.

In summary, inventories play two major roles in our model. On the firm's side, inventories allow the manufacturer to engage in cost-minimizing production bunching. On the consumer's side, higher levels of inventories provide more variety, thus making it easier to match consumers with their ideal vehicle. Hence, pricing and inventory decisions are linked both through the firm's cost structure as well as the demand system for automobiles.

Our two-sided model provides a consistent explanation of the four stylized facts. We replicate facts (i) and (iii) by modelling the firm as solving an inventory control problem while facing declining demand over the product cycle. Early in the model year, the automaker sets prices sufficiently high to keep sales less than production to accumulate a large stock of inventories. Building up inventories, or following a build-to-stock inventory management strategy, is optimal because it strengthens demand by increasing variety. Over the remainder of the model year, our estimate of leftward-shifting demand lowers the shadow value of inventories (i.e., the marginal cost curve), resulting in a 9.0% decline in the price over the entire product cycle and an average vintage premium of 8.5% (fact (ii)). Because inventories are used to both optimally schedule production and increase variety, the model is able to match the high level of inventories relative to sales (fact (iv)).

An innovation in this article is to explicitly model how inventories can bolster demand by increasing the variety of vehicles available to consumers. To quantify the importance of this role for inventories, we simulate the model under a counterfactual build-to-order strategy. When a manufacturer builds automobiles according to orders, the firm is able to offer consumers full variety for every product without holding inventories. Hence, the role for inventories in bolstering demand is shut down. Under this alternative policy, we find that automakers' pricing strategies are significantly different: within-model-year prices decline by 5.3%, roughly six tenths of the percent price decline observed under the firms' current build-to-stock policy. A main result,

¹ Federgruen and Heching (1999) and Elmaghraby and Keskinocak (2003) provide a nice overview of the more recent revenue management literature within operations research.

then, is the model's prediction that automakers' build-to-stock inventory management policy is responsible for four tenths of the 9.0% decline (annual rate) in prices over the model year. More generally, our work suggests a significant driver behind a durable good's price decline may be the firm's inventory management strategy.²

Our work builds on the macroinventory literature. Typically, this literature recognizes a role for inventories in spurring demand, but then assumes an exogenously specified target inventory-sales ratio into the firm's problem (e.g., Blanchard, 1983).³ In contrast, our work puts more structure on the effect of inventories on demand and provides an estimate of the elasticity between unit sales and inventories in stock. Significantly, the optimal inventory-to-sales ratio is endogenous in our model.

In addition, our work builds on the microinventory literature. Work by Reagan (1982), Aguirregabiria (1999), and Zettelmeyer, Scott Morton, and Silva-Risso (2006), for example, study the interactions between pricing and inventory management. Our work is closest in spirit to Aguirregabiria (1999), who estimates a structural model of a retailer which accounts for the joint dynamics of prices, sales, and firm-held inventories. His work demonstrates that the possibility of stockouts along with fixed ordering costs can explain the high-low pricing schemes retailers often employ. In contrast, we focus on a durable goods market and consider inventory's role in increasing the variety available to consumers. We demonstrate how this demand for inventories partly explains the observed price decline of vehicles over the product cycle.

2. Data sources and empirical observations

■ In this section, we outline our data sources and document four stylized facts.

□ **Data sources.** To construct a data set of transaction prices, sales, production, and inventories by model and model year in the United States, we combined data from two sources. The first data source includes detailed information on U.S. retail transactions collected from a sample of vehicle dealerships. It reports prices, by model and model year, and the distribution of model-level sales across model years. The second data source reports total sales in North America, by country and model, and on production, by model and model year.

The first data set was constructed by Corrado, Dunn, and Otoo (forthcoming), using data from J.D. Power and Associates (JDPA). JDPA collects daily transaction-level information from dealerships across the United States. JDPA aggregated these data to generate a monthly time series of average price, sales, average cash rebate, and average financial package by model and model year (e.g., 2000 Ford Escort). Our sample covers the period from January 1999 to January 2004 and represents 70% of the geographical markets in the United States and roughly 15%–20% of national retail transactions. JDPA attempts to precisely measure the transaction price of a vehicle. The price they obtain includes the price of accessories (such as roof racks) and transportation costs but excludes aftermarket options, taxes, title fees, and other document preparation costs. Further, JDPA adjusts this price to account for instances when a dealership undervalues or overvalues a customer's trade-in vehicle as part of a new vehicle sale.⁴ JDPA's transaction price does not account for incentives the customer received to help finance the purchase of the car; hence, we define the average market price of a vehicle as the transaction price minus the cash rebate minus a measure of the financial incentive offered by the manufacturer.

Our measure of the value of the financial incentive is based on the amount financed, interest rate, and loan term that the average customer received. JDPA captures these financial data when

² Durable goods price declines over the product cycle are not unique to the automobile industry, having been documented for a number of other products, including textbooks (Chevalier and Goolsbee, 2009), microprocessors (Aizcorbe and Kortum, 2005), and consumer electronics (Gowrisankaran and Rysman, 2010; Copeland and Shapiro, 2010).

³ Bills and Kahn (2000) provide a good synopsis of the different ways the inventory literature has built in a demand for inventories outside of its production-smoothing role.

⁴ A trade-in vehicle's benchmark value is the wholesale price.

loans are arranged through the dealership. As a majority of car loans arranged through dealerships are made by the financing arms of manufacturers, we treat the financial data as an approximation of the average financial package that consumers received from manufacturers. To measure the value of these financial incentives to consumers, we compare the financial package in the data against a benchmark package offered by commercial banks. We make this comparison by first computing the net present value (NPV) of the average amount financed given the interest rate and loan term in the data. We then compute the NPV of financing the same average amount at the average interest rate reported for 48 month new car loans at commercial banks. The value of the manufacturer's financial incentive is then defined as the difference between the two NPV amounts, deflated by the Bureau of Economic Analysis's (BEA's) personal consumption deflator.

We linked the JDPa data set to a data set from Ward's Communications on the U.S. sales and North American production of General Motors, Ford, and Chrysler (a.k.a. the Big Three). We excluded foreign manufacturers because measuring overseas production is difficult. The sales data for the Big Three are available only at the model level, not by model year. Therefore, we constructed estimates of sales by model and model year using the distributions of sales across model years in the JDPa sample. Using model changeover dates at assembly plants, we decomposed the production data by model into observations by model year. Finally, using the sales and production estimates by model and model year, we constructed estimates of inventories over the sample period. All told, the work described here results in a data set with monthly observations, by model year, on the real average price, quantity sold, quantity produced, and inventory held for almost all light vehicle models sold by the Big Three in the United States from 1999 to 2003.

□ **Empirical observations.** As described in the Introduction, we observe several stylized facts that hold across models and model years. To provide an illustrative example, we plot in Figures 1 the price, sales, production, and inventory data for a typical midsize car. In Figure 1A, We see a steady decrease in price for each vintage. In the 2000 model year, the average price for the midsize car falls over \$2000, more than 10% of the initial price. The declines in prices for subsequent model years are just as pronounced. Figures 1A and 1B exhibit the simultaneous sale of multiple vintages as well as the premium the newer model-year vehicle commands over the older model-year vehicle. We refer to this difference in price as the "new vintage premium." The size of this premium varies, but the average premium for this particular midsize car is 7%. In Figures 1B and 1D, we see that sales and inventories exhibit hump-shaped profiles. Finally, Figure 1C illustrates the large volatility in vehicle production, a consequence of frequent weeklong plant shutdowns.

These patterns hold at the aggregate level. To observe the within-year price declines more generally (fact (i)), Figure 2 illustrates the aggregate matched-model price indexes for successive model years constructed by Corrado, Dunn, and Otoo (forthcoming) using the entire JDPa data set.⁵ As can be seen, transaction prices for a given model year are highest at the model's introduction and trend downward over the course of the product cycle. Table 1 provides a summary of the average monthly price decline by market segment and model year. For the midsize market segment, the mean monthly price decline of 1999 model-year vehicles is 9.1% at an annual rate. On average, midsize automobiles fall 9.2%. Table 1 illustrates the wide range in average price declines both across market segments and model years.⁶ In general, luxury vehicles decline the most in price, followed by pickup trucks. Looking across model years, 2003 vehicles decline the most in price by far, reflecting especially high incentives offered by manufacturers in the latter half of the product cycle. Overall, the monthly decline in price averages 9.0% at an annual rate.

The overlap of the model-year price indexes highlights the second stylized fact: multiple vintages of a model are sold simultaneously. This is accomplished by selling the older vintage out

⁵ The price declines illustrated in Figure 2 will be smaller than those reported elsewhere in this article, because they include European and Asian manufacturers. The Big Three were the most aggressive in cutting prices over the model year.

⁶ We exclude the van market segment from our analysis because a substantial number of vans are sold to firms.

FIGURE 1

PRICES, SALES, PRODUCTION, AND INVENTORIES FOR A TYPICAL MIDSIZE CAR BY MODEL YEAR

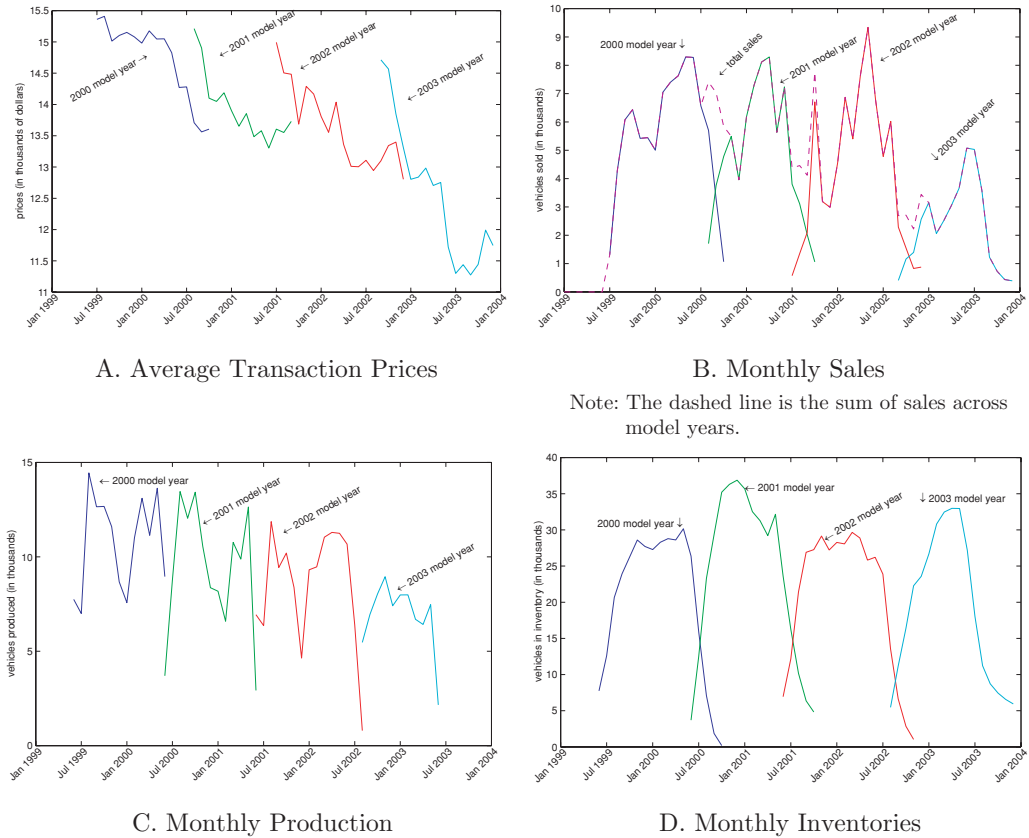


TABLE 1 The Monthly Price Decline (Annual Rate) by Market Segment and Model Year


Market Segment	Model Year					Average	
	1999	2000	2001	2002	2003		
Compact	7.7	5.9	8.1	9.4	17.5	9.5	(2.4)
Midsize	9.1	6.7	6.2	8.9	16.3	9.2	(1.5)
Full-size	8.9	7.9	6.4	8.5	13.4	8.9	(2.1)
Luxury	12.2	11.2	9.3	13.3	15.3	12.1	(1.4)
Pickup	6.7	9.5	5.3	8.6	16.7	9.6	(2.2)
SUV	7.0	6.7	7.1	5.2	13.6	8.2	(0.9)
Sporty	2.1	6.2	0.2	6.0	10.9	4.9	(2.5)
Average	7.7	7.6	6.4	7.9	15.4	9.0	(0.7)

Note: Standard errors are in parentheses. Cells are the sales-weighted average of the monthly percentage change in price for all Big Three light vehicles sold.

of inventories. In our sample, the typical vehicle is produced for 12 months, but is on the market for 16.7 months. Hence, automakers find it profitable to substantially extend a model's life and so sell two vintages of the same model simultaneously. The number of months sold varies little across types of vehicles; the mean length of the automobile product cycle has a standard error of only 0.02.

FIGURE 2

MATCHED-MODEL PRICE INDEXES BY MODEL YEAR



Source: Figure taken from Corrado, Dunn, and Otoo (forthcoming). It was constructed using price data on all vehicles sold in the United States.

TABLE 2 The Average New-Vintage Premium by Market Segment and Model Year

Market Segment	Model Year					Average	
	2000	2001	2002	2003	2004		
Compact	6.5	6.9	6.9	7.5	12.5	7.8	(0.5)
Midsize	8.7	5.4	6.1	6.7	12.4	8.1	(0.4)
Full-size	9.6	7.1	8.4	8.8	8.8	8.5	(0.5)
Luxury	13.1	10.8	9.3	15.9	10.6	12.0	(0.4)
Pickup	11.2	8.8	7.1	10.5	23.4	12.0	(0.9)
SUV	5.1	0.8	9.8	8.8	10.9	7.3	(0.3)
Sporty	2.3	7.6	3.5	28.1	-9.4	6.4	(0.8)
Average	8.2	5.5	7.6	9.4	14.0	9.0	(0.2)

Note: Standard errors are in parentheses. Cells are the sales-weighted average of the percentage difference in price between the new and old vintage of the same model, conditional on both vintages being sold in the same month. The sample includes all Big Three light vehicle sales.

The combination of decreasing prices over the model year and the simultaneous sale of multiple vintages implies that newer vintages command a premium over their older counterparts. In Table 2, we report the average new-vintage premium by market segment and model year. The new-vintage premium varies quite a bit across market segments and model years, with an overall average of 9.0%. Across model years, the average new-vintage premium is typically between 5%–9%, although the premium during the 2003 to 2004 changeover is 14.0%. This large premium is related to the steep decline in prices for 2003 model-year vehicles, as shown in Table 1.

One might argue that the new-vintage premium simply reflects improvements in quality or additional features. For example, the 23.4% new-vintage premium recorded for 2004 model-year pickup trucks reflects, in part, a quality improvement made to Ford's F-series. Alternatively, if a cheaper base model is introduced, as was the case with the 2004 Ford Mustang, the vehicle premium may be biased downward; note the -9.4 premium for 2004 sporty cars. For the large majority of the vehicles in our sample, however, changes in the observable characteristics from one model year to the next were minimal, and even for vehicles with such changes, the downward-sloping price pattern was still apparent. To further investigate, we recomputed the new-vintage premia excluding vehicles that had undergone a major redesign and found that these new premia differed little from the magnitudes reported in Table 2.

The decline in an automobile's price over the model year and the resulting new-vintage premium has been studied by Pashigian, Bowen, and Gould (1995). They hypothesize that prices decline because the fashion component of a vehicle depreciates. Given our data, we posit that within-model-year price declines are driven more by the used-vehicle market than by fashion. Used-vehicle prices are mainly a function of their model year, not their date of production. Hence, even if a 2001 and a 2000 model-year vehicle of the same model are produced just days apart and are similar in observed characteristics, their value on the used-car market are substantially different.

To provide evidence in support of this hypothesis, we estimate a price regression on a separate JDPa data set of *used-vehicle* transactions from 2001 to 2003. The left-hand-side variable is the log of the transaction price for a given model and vintage of a used vehicle. On the right-hand side, we measure a vehicle's age by the calendar year minus the model year plus one. Because physical depreciation can significantly influence price, as a proxy for wear and tear, we include the vehicle's odometer reading when sold. Finally, we also include calendar month and year dummies, model fixed effects, and vehicle characteristics such as engine size to capture differences in price not due to a vehicle's age. This regression then, should capture changes in price due to age, controlling for a host of observable characteristics. To further control for changes in quality across vintages of the same model, we restrict the sample to vehicles of age 4 or under. This restriction reduces the variation in price across vintages of the same model due to changes in unobserved characteristics. With these restrictions, we have 34,685 observations. The regression fits the data well; the R^2 is 0.987. The estimated coefficient on age is -0.097 and is statistically significant. It implies that, even after controlling for the odometer reading and other vehicle characteristics, increasing model age by 1 year decreases the value of a used vehicle by 9.7%, a figure only slightly greater than our estimate of the new-vintage premium. In line with our expectations, the coefficient on odometer is also statistically significant and implies a price decline of about 0.4% for each additional 1000 miles driven. These results point to the importance of a vehicle's age in the used-car market, and strongly suggest that the new-vintage premium is partly driven by the difference in the new vehicles' values in the used-vehicle market.

Turning to fact (iii), Figures 3 and 4 show the slow rise of aggregate sales and inventories over the first 6 months of the model year. Both time series then plateau for several months before falling off over the tail end of the model year. To better analyze the relationship between sales and inventories, we consider the ratio of inventories to sales, also known as days-supply. This ratio measures the number of days the firm could continue to sell cars if it used only the stock of inventories available at the start of the month. On average, automakers carry 75 days-supply (fact (iv)), or enough inventories to sell vehicles for over 3 months without any additional production! Table 3 provides a breakdown of the average days-supply by market segment and illustrates the substantial variation in days-supply across different types of vehicles.

For the remainder of the article, we present a model designed to replicate these empirical regularities. We first describe the firm's problem. We assume the automaker takes market demand curves as given and solves a dynamic profit maximization problem. As the automaker is able to hold inventories, at certain times it is able to sell two vehicles, the current year's vintage and the previous year's vintage. We posit a log-log market demand curve whose parameters are

FIGURE 3

BIG THREE AGGREGATE U.S. SALES BY MODEL YEAR




FIGURE 4

BIG THREE AGGREGATE U.S. INVENTORIES BY MODEL YEAR




TABLE 3 The Average Days-Supply by Market Segment

Market Segment	Compact	Midsize	Full-size	Luxury	Pickup	SUV	Sporty	Average
Days-Supply	73	65	75	80	84	75	83	75 (2.4)

Note: Standard errors are in parentheses. Days-supply is the ratio of inventory to sales, and states the number of days the current flow of sales could be sustained if the only source of vehicles is the current stock of inventories. To convert this measure from months to weeks to days, we multiplied the ratio by 4.3 and 6. The sample includes all Big Three light vehicle sales.

elasticities with respect to prices and product variety. We then draw upon the existing discrete-choice literature to estimate these elasticities. The supply-side parameters of the firm's problem are chosen to match the key features of the firm's cost structure and the means of prices and inventories. We derive decision rules that govern the production and pricing of vehicles over the model year. Through numerical simulations, we demonstrate that our derived decision rules under a build-to-stock inventory policy are consistent with these stylized facts.

3. An industry model with overlapping vintages

■ In the interest of tractability, we make several assumptions on the supply side. First, each vehicle line within the firm can be considered a separate, independent subfirm or profit center. Hence, an automaker is modelled as a collection of dynamic programs that can be solved independently of each other. Second, we integrate the dealership into the automaker and consider a unified pricing decision. Third, we abstract from bargaining by assuming that all customers who purchase during a particular period pay the same retail price. Of course, there are many interesting questions about how the automakers decentralize their operations both across products and between the production and marketing sides of the business.⁷ Although issues of vertical control and price discrimination are present in the automobile market, we are implicitly assuming that manufacturers and dealers jointly set prices to maximize their combined profits and solve the double-marginalization problem.⁸ Furthermore, we interpret high levels of inventories nationally to reflect high levels of inventories at all dealerships. Hence, automakers are able to coordinate with dealerships so that there is not an uneven distribution of inventories across the country.⁹

The automaker produces one vintage of a vehicle at a time, switching to build a newer vintage every year. Although production is exogenously limited to 52 weeks, the number of weeks a vehicle is sold is endogenous. In particular, through the use of inventories, the automaker can sell a specific vintage for more than a year. This also implies that an automaker can choose to sell two vintages of a vehicle at the same time. A specific vintage is labeled this year's vintage, or the new-vintage, for the first 52 weeks of its life. After 52 weeks, when it is no longer being produced, we label the specific vintage last year's, or the old, vintage. The automaker's decision period is a week, where the automaker solves an infinite horizon problem by repeatedly solving a 52 week problem. Each week the firm must decide (i) the number of vehicles of the current model year to produce, q_i ; (ii) the number of days to operate the plant, D_i , the number of shifts

⁷ For example, Bresnahan and Reiss (1985) model and estimate the division of markups between automobile manufacturers and dealers. Busse, Silva-Risso, and Zettelmeyer (2006) estimate how the value of manufacturers' incentive programs are split between dealers and final customers.

⁸ For discussions of bargaining and price discrimination in the retail auto market, see Ayres and Siegelman (1995), Goldberg (1996), Scott Morton, Zettelmeyer, and Silva-Risso (2003), and Langer (2009).

⁹ Dealerships of the same brand might compete with one another through inventory accumulation. This seems unlikely, however, because dealerships of the same brand are strategically located to minimize within-brand competition. Further, our understanding of the industry is that manufacturers distribute vehicles with an eye toward minimizing strategic games among dealerships.

to run, S_t , and the number of hours per shift, h_t ; (iii) the retail price of the current vintage, p_t^{this} ; and (iv) the retail price of last year's vintage, p_t^{last} (if any are still in stock).

We assume that weekly sales, s_t^j , for each of the two vintages depend on each vintage's own price, the price of the other vintage, and the stock of each vintage, I_t^j , that is inventoried at the end of period $t - 1$:

$$\log s_t^j = \mu_t^j - \eta_t^j \log(p_t^j) + \phi_t^{ji} \log(p_t^i) + \zeta_t^j \log\left(\frac{I_t^j}{I_{mean}^j}\right) \quad \text{for } j, i = \{\text{this, last}\} \text{ and } i \neq j, \quad (1)$$

where μ_t^j is a constant term, η_t^j is the own-price elasticity, ϕ_t^{ji} is the cross-price elasticity, and ζ_t^j is the own-variety elasticity. These elasticities may vary across the 52 weeks of the year. With the variety term, $\frac{I_t^j}{I_{mean}^j}$, we seek to capture the idea that consumers are more likely to purchase a vehicle if they can find one that matches their particular tastes. Within the automobile industry, variety means having vehicles on a dealership lot with all possible combinations of options (e.g., color, leather interior, automatic transmission). Hence, our definition of variety translates into a measure of the number of vehicles at a dealership. Because we do not have data at the dealership level, our proxy for variety is inventories (i.e., the number of cars at all dealerships) divided by the mean level of inventories for the appropriate market segment. We do not simply use the level of inventories as our measure of variety because the number of dealerships by market segment varies. Intuitively, vehicles that appeal to buyers across the United States will require larger amounts of inventory to achieve the same level of variety, relative to less popular vehicles only sold in parts of the country. Mercedes-Benz, for example, only had 191 dealerships in the United States in 2002, whereas Honda had 959.¹⁰ Dividing through by the mean allows us to compare the inventory accumulation of popular vehicles such as pickups, and its resulting effect on variety, to other vehicles. A natural question is why we did not use a more disaggregate mean level of inventories. After all, even within market segments there is variation in vehicle popularity. Given that our model/model-year inventory measures are inferred from estimated sales and production flows, however, we are worried about the level of noise in the data at this level. Further, we would need to impute mean inventory levels for a number of models for which we do not observe a full model year (e.g., models newly introduced at the end of the sample).

Our inventory-based measure of variety assumes that higher levels of inventory imply higher levels of variety. Unfortunately, we do not have any direct evidence this is true nor do we have alternative measures of variety. Nevertheless, linking higher levels of inventory with more variety in an industry with significant product differentiation seems reasonable and is consistent with results reported in Cachon and Olivares (2010).

Because there is no intercept with constant-elasticity demand curves, we assume that customers never pay more for last year's vintage than for the current vintage:

$$s_t^{last} = 0 \quad \text{if } p_t^{last} > p_t^{this}. \quad (2)$$

Unsold vehicles can be inventoried without depreciation. Current production is not available for immediate sale, so sales can be made only from the beginning-of-period inventories:

$$s_t^j \leq I_t^j. \quad (3)$$

Further, sales cannot be backlogged. Inventories for the current vintage follow the standard law of motion:

$$I_{t+1}^{this} = I_t^{this} + q_t - s_t^{this}. \quad (4)$$

Because no vehicles for the last model year are produced during the current year, inventories for last year's vintage evolve according to

$$I_{t+1}^{last} = I_t^{last} - s_t^{last}. \quad (5)$$

¹⁰ Data taken from Ward's 2002 *Automotive Yearbook*.

At the conclusion of the current model year, any unsold vehicles of last year’s vintage are scrapped at a zero price, and this year’s vintage becomes last year’s vintage:

$$I_1^{last} = I_{s2}^{this} + q_{s2} - s_{s2}^{this} . \tag{6}$$

We assume the vehicle is assembled at a single plant. Each period, the firm must decide how many vehicles of the current vintage to produce and how to organize production to minimize costs. As documented by Hamermesh (1989) and Bresnahan and Ramey (1994), managers at durable goods manufacturing plants typically increase or decrease production by altering the length of the work week rather than the rate of production (i.e., the speed of the assembly line). In particular, weeklong plant shutdowns are frequently employed. In auto industry lingo, these plant closures are called inventory adjustment shutdowns. In order to induce the firm in our model to engage in similar production scheduling, we assume the firm has a linear production function but faces a set of nonconvex costs.

We assume the plant can operate D days a week. It can run one or two shifts, S , each day, and both shifts are h hours long. We assume the number of employees per shift, n , and the line speed, LS , are fixed. So the firm’s production function is

$$q_t = D_t \times S_t \times h_t \times LS. \tag{7}$$

We assume the firm faces a set of nonconvex costs to running the plant each week. We motivate these nonconvex costs from the union contract, although we recognize that the contract structure is endogenous and that the nonconvexities may be due to the underlying technology.

From the autoworkers’ union contracts, we know that workers on the second shift receive a 5% premium above the first-shift wage. Any work in excess of 8 hours a day, and all Saturday work, are paid at a rate of time and a half. Employees who work fewer than 40 hours per week must be paid 85% of their hourly wage times the difference between 40 and the number of hours worked. This “short-week compensation” is in addition to the wages paid for hours actually worked. If the firm chooses to not operate a plant for a week, the workers are laid off. Laid-off workers receive 95 cents on the dollar of their 40 hour pay in unemployment compensation. Of these 95 cents, the firm pays about 65 cents.

Given such a labor contract, if the firm decides to produce q vehicles, it must then choose how many days to operate the plant, how many shifts to run, and how many hours to run each shift to minimize its cost of production. Given these choices, the firm’s week t cost function is expressed as

$$c(D_t, S_t, h_t | q_t) = \gamma q_t + (w_1 + I(S_t = 2)w_2) \times (D_t h_t n + \max[0, 0.85(40 - D_t h_t)n] + \max[0, 0.5D_t(h_t - 8)n] + \max[0, 0.5(D_t - 5)8n]) + 0.65w_1 40(2 - S_t)n, \tag{8}$$

where n is the number of employees per shift, and w_1 and w_2 are the hourly wage rates paid to the first-shift and second-shift workers, respectively. γ is the per-vehicle material cost and incorporates all costs (such as materials, energy, transaction) that do not depend on the allocation of production over the week. The first term within the brackets represents the straight-time wages paid to the production workers, and the subsequent terms capture the 85% short-week rule and the overtime premia. The last term is the unemployment compensation bill charged to the firm. This cost function is piecewise linear with kinks at one shift running 40 hours per week and two shifts running 40 hours per week. This implies that the firm minimizes average cost by operating the plant with either one shift or two shifts for 40 hours per week. When the plant operates below its minimum efficient scale, the cost-minimizing production schedule involves bunching production

by oscillating between running two 40 hour shifts for several weeks and then shutting down the plant for a week.¹¹

The firm’s objective is to maximize the present value of the discounted stream of profits. For each model year, the automaker’s problem is to maximize

$$\sum_{t=1}^{52} \left(\frac{1}{1+r} \right)^{t-1} \{ p_t^{last} s_t^{last} + p_t^{this} s_t^{this} - c(D_t, S_t, h_t | q_t) \} + \left(\frac{1}{1+r} \right)^{52} V(I_1^{last}, 0, 1) \quad (9)$$

subject to (1)–(7) and where $c(D, S, h|q)$ is given by (8). The term $V(I_1^{last}, 0, 1)$ is a continuation value, which we now define.

Let $V(I^{last}, I^{this}, t)$ be the optimal value at week t for the firm that holds in inventory I^{last} of last year’s vintage and I^{this} of this year’s vintage. Then the firm’s value function for $t = 1, 2, \dots, 51$ can be written:

$$V(I^{last}, I^{this}, t) = \max_{p^{this}, p^{last}, q} \left\{ p^{last} s^{last} + p^{this} s^{this} - \min_{D,S,h} c(D, S, h | q) + \frac{1}{1+r} V(I^{last} - s^{last}, I^{this} + q - s^{this}, t + 1) \right\} \quad (10)$$

subject to (1), (2), (3), and (7) and where $c(D, S, h|q)$ is given by (8). At week 52, this year’s vintage becomes last year’s vintage, and so the value function is

$$V(I^{last}, I^{this}, 52) = \max_{p^{this}, p^{last}, q} \left\{ p^{last} s^{last} + p^{this} s^{this} - \min_{D,S,h} c(D, S, h | q) + \frac{1}{1+r} V(I^{this} + q - s^{this}, 0, 1) \right\}. \quad (11)$$

Following a suggestion by John Rust, we merge the 52 value functions into a single time-invariant Bellman equation:

$$V(I^{last}, 0, 1) = \max_{\{p_t^{this}, p_t^{last}, q_t, D_t, S_t, h_t\}} \left\{ \sum_{t=1}^{52} \left(\frac{1}{1+r} \right)^{t-1} (p_t^{last} s_t^{last} + p_t^{this} s_t^{this} - c(D_t, S_t, h_t | q_t)) + \left(\frac{1}{1+r} \right)^{52} V(I_{52}^{this} + q_{52} - s_{52}^{this}, 0, 1) \right\}. \quad (12)$$

For a given parameter vector, we carried out the following steps to solve for the fixed point: (i) guess an initial value for $V(I^{last}, 0, 1)$; (ii) solve the 52 Bellman equations in (10) and (11) through backward recursion; (iii) compute a new value for $V(I^{last}, 0, 1)$ through policy iteration; and (iv) repeat steps 2 and 3 until a fixed point is reached. More details on the solution method are provided in Appendix B.

4. Parameterizing the model

■ There are a large number of parameters in this model. For the demand-side parameters, we employ a discrete-choice methodology to estimate consumers’ preferences over automobiles. We then use these estimates to compute the intercepts, own-price elasticities, cross-price elasticities, and own-variety elasticities that are parameters in the market demand function, equation (1). For the supply-side parameters, we choose some values based on published information on assembly plants. The remaining values are set to match a set of first moments in the data.

¹¹ If we assumed a convex cost function, the main results of this article would still go through. We incorporate this nonconvex cost structure because one reason automobile firms hold inventories is to facilitate plant shutdowns due to scheduled holidays or a desire to reduce production.

□ **Demand-side parameters.**

Overview. Following Berry, Levinsohn, and Pakes (1995), henceforth BLP, we construct the demand system by aggregating over the discrete choices of heterogeneous individuals. The utility derived from choosing an automobile depends on the interaction between a consumer's characteristics and a product's characteristics. Consumers are heterogeneous in income as well as in their tastes for certain product characteristics. We distinguish between two types of product characteristics: those that are observed by the econometrician (such as size and horsepower), which are denoted by X ; and those that are unobserved by the econometrician (such as styling or prestige), which are denoted by ξ . We allow for households' distaste for price, denoted by α , to vary from quarter to quarter, capturing the possibility that different types of households show up to purchase a new automobile at different times of the year. We specify the indirect utility derived from consumer i purchasing product j in period t as

$$u_{ijt} = X_{jt}\beta + \xi_{jt} - \sum_{q=1}^4 1_{d_t=q} \alpha_{iq} p_{jt} + \sum_k \sigma_k v_{ik} x_{jkt} + \varepsilon_{ijt}, \quad (13)$$

where p_{jt} denotes the price of product j in period t and $x_{jkt} \in X_j$ is the k th observable characteristic of product j . Let d_t denote the quarter of the automotive year into which period t falls, and let $1_{d_t=q}$ be an indicator variable equal to 1 when d_t is equal to $q \in \{1, 2, 3, 4\}$. The term $X_{jt}\beta + \xi_{jt}$, where β are parameters to be estimated, represents the utility from product j that is common to all consumers, or a mean level of utility, δ_{jt} . Consumers then have a distribution of tastes for each observable characteristic. For each characteristic k , consumer i has a taste v_{ik} , which is drawn from an independently and identically distributed (i.i.d.) standard normal distribution. The parameter σ_k captures the variance in consumer tastes. The term α_{iq} measures a consumer's distaste for price. Following Berry, Levinsohn, and Pakes (1999), we assume that $\alpha_{iq} = \frac{\alpha_q}{y_i}$, where α_q is a parameter to be estimated and y_i is a draw from the income distribution. Finally, ε_{ijt} is distributed i.i.d. type-1 extreme value.

Consumers choose among the $j = 1, 2, \dots, J$ automobiles in our sample and the outside good, which represents the choice not to buy a new automobile from the Big Three. Consumers maximize utility, and market shares are obtained by aggregating over consumers.

Implementation. As described in Section 2, our sample includes data on the Big Three firms over the 5 year period from February 1999 to January 2004. There are 638 observations of unique model and model-year vehicles. Industry wisdom is that consumers sometimes time their vehicle purchase decisions, for example to take advantage of end-of-month sales. As such, we believe our static demand model is better suited to analyzing quarterly, rather than monthly, data. Hence, we aggregate sales and prices to the quarterly frequency.

As was done in previous research, we link sales and prices to the characteristics of the base model to produce a vehicle-quarter observation.¹² Following Nevo (2001), we use model-level fixed effects as the matrix of observable characteristics used to compute the mean utility of a product. We supplement these dummies with a quadratic time trend, model-year dummies, and measures of congestion, variety, and "newness." The congestion dummy variable draws from the work of Akerberg and Rysman (2005), who demonstrate the importance of controlling for variation in the choice set when estimating consumers' price elasticities. Because of the overlap in model years, households face large variation in the number of products offered over time. To capture this effect, we use an indicator variable equal to 1 when two vintages of the same model are sold in the same quarter. The variety term, to our knowledge, has not been incorporated into the BLP framework. We use the definition of variety described above (see Section 3), the ratio of inventories to the mean level of inventories for the appropriate market segment. To better capture the substitution patterns between two vintages of the same model, we use a newness dummy variable equal to one if a model has been sold for less than a year.

¹² Information on vehicle characteristics were taken from Automotive News's *Market Data Book* (various years).

Finally, measures of acceleration and dimensions, along with the newness dummy variable and constant term, are included in the vector of observable characteristics used to measure heterogeneity in households' preferences, $\sum_k \sigma_k v_{ik} x_{jkt}$.

Following BLP, we use the number of households in the United States as reported in the Current Population Survey (CPS) as a measure of market size for the year. Because we do not have any information on the number of households who are actively shopping for automobiles throughout the year, we assume that one fourth of all households in a given year show up each quarter.¹³ We assume the distribution of household income is lognormal, and, for each year in our sample, we estimate its mean and variance from the CPS.

Our estimation strategy follows the generalized method of moments approach taken by BLP.¹⁴ We match the usual moments, that the expected value of ξ , conditional on the observed characteristics, is equal to zero, $E[\xi | X] = 0$. Because ξ is correlated with price, an endogeneity problem arises.¹⁵ We follow BLP and use competing products' characteristics as instruments.

Although characteristics only vary at the model-year frequency, the overlap of different vintages along with some timing differences in the introduction of new vehicles over the calendar year provides enough variation at the quarterly frequency. To demonstrate the impact of our instruments, we run a simple logit version of our demand model with and without instruments.¹⁶ The noninstrumented estimate of price is -0.232 , whereas the instrumented price estimate is -0.367 , were both estimates where highly significant. Our instruments, then, do have a substantial impact on the estimated price coefficient. The level of inventories could plausibly be considered endogenous as well. The noninstrumented and instrumented estimates of the coefficient on the variety term, however, are similar.

Because we include an inventory-based variety term in the demand estimation, our moment conditions assume that the current level of beginning-of-period inventories are orthogonal to ξ_{jt} . Because this period's inventories are a function of $\xi_{j,t-1}$, our moment conditions rule out serial correlation in ξ_{jt} (after controlling for model-level fixed effects), a relatively strong assumption. If this assumption is incorrect, we think the most likely outcome would be demand residuals that are positively correlated over time. This would lead to inventories being negatively correlated with the current demand shock, suggesting that our estimate of the coefficient on inventories is downward biased.

Our model follows the same identification strategy as BLP and the accompanying literature. Consumers' sensitivity to price is identified by observing changes in market share alongside changes in prices. Further, consumers' sensitivity to price is identified through changes in market share alongside changes in the choice set. As detailed in Section 2, we observe substantial variation in vehicle prices. In addition, there are changes in the consumers' choice set because automakers introduce new vehicle vintages while continuing to sell old vintages.¹⁷ The elasticity of sales with respect to inventories is similarly identified. The time series of inventories for a particular vehicle is different across model years because automakers will under- and overpredict demand. Hence, through the time series dimension of the data, our model can determine whether having a large stock of inventories for a particular vehicle, all else equal, generates greater sales. The cross-section also provides identification. Because the firm-held stock of inventories for each

¹³ We tried an alternative approach that links the number of households per quarter to total light motor vehicle sales, generating a correlation between the outside good's share and the number of households. The estimated parameters and implied elasticities, however, did not significantly change with this alternative definition of market size.

¹⁴ We modified the programs provided in Nevo (2000) to estimate the demand system. A notable addition to this set of programs is the importance sampling simulator described in BLP, used to reduce sampling error.

¹⁵ Berry (1994) provides a detailed explanation of, and solution to, this problem for discrete-choice demand models.

¹⁶ For this logit model, the dependent variable is the log difference of a product's and outside option's market share. The independent variables are those in the linear portion of consumers' indirect utility for our demand-side model.

¹⁷ As detailed in Akerberg and Rysman (2005), large changes to the consumer's choice set could be a cause of concern with respect to identification of the price elasticity. Following their recommendation, we address this problem by controlling for congestion in the product space.

TABLE 4 Parameter Estimates

Parameters		Coefficient	Standard Error
Heterogeneity in Tastes	σ		
	Constant	2.50	1.23
	Acceleration	0.14	0.51
	Height	0.92	1.61
	Size	0.58	0.44
	Newness	1.82	0.26
Variety		0.58	0.10
Congestion		-0.72	0.08
Newness		0.94	0.30
Distaste for price (Q1)	α_1	25.78	2.32
Distaste for price (Q2)	α_2	26.04	5.12
Distaste for price (Q3)	α_3	23.09	4.96
Distaste for price (Q4)	α_4	19.49	5.01

Note: Q1, Q2, Q3 and Q4 denote the first through fourth quarters of the automotive year, which starts in August. Standard errors have been corrected for serial correlation of the unobserved characteristic within a vehicle across quarters.

vehicle will differ, the model identifies how differences in inventory explain differences in market shares, controlling for price and other characteristics.

Results. We present a subset of the parameter estimates in Table 4. Given their large number, we do not report all our estimates on the linear portion of utility (β in equation (13)). Instead, we show the estimates of the congestion, variety, and newness coefficients, the consumers' distaste for price (α), and the measure of the heterogeneity in consumers' tastes (σ). The standard errors reported in Table 4 have been corrected for serial correlation of ξ within a vehicle (i.e., a given model/model year) across quarters.

The coefficients on the acceleration, height, and size are not statistically significant. However, we estimate that consumers are quite heterogeneous in their tastes for purchasing a new car (i.e., the constant term) and in their tastes for purchasing a new car at the beginning of the model year (i.e., the newness variable). The positive, significant value of the variety term accords with our prior belief that more variety is valued by consumers. The magnitude of this coefficient is more easily appreciated in terms of an elasticity, which is discussed below. The negative and significant value of the congestion term indicates that congestion is important in the automobile market when considering the overlap in vintages. As detailed in Ackerberg and Rysman (2005), this result shows the importance for flexibility in the i.i.d. logit errors across different vintages of the same model. Otherwise, the estimated parameters (especially the price coefficients) could be biased.

Most importantly, the price coefficients are precisely estimated. The estimated value of households' distaste for price is in the neighborhood of 25, although there is a dropoff in its value in the fourth quarter. The quarters differ from calendar quarters. We defined the first quarter as the first 3 months of a typical vehicle's product cycle: August, September, and October. We then defined the second through fourth quarters on the basis of this new grouping of months.

The magnitude of the variety and price coefficients are more easily interpreted by examining the appropriate elasticities. We start with the most important set of elasticities, the own-price elasticities (Table 5). We report the average of individual elasticities across market segments, quarters, and vintages, where the vintage label signifies whether the vehicle is the newest model year available or not.

The own-price elasticities generated by our parameter estimates range between 1.5 and 3, indicating that manufacturers face elastic demand. In the first quarter a car is sold, our results imply that a 1% price increase for a typical compact car (roughly \$140) causes a 2.6% fall in sales, holding everything else equal. Own-price elasticities vary little across quarters. In general, our

TABLE 5 The Average Absolute Value of Own-Price Elasticities by Market Segment, Quarter, and Vintage

Vintage	Market Segment	1st Quarter	2d Quarter	3d Quarter	4th Quarter
New	Compact	2.6	1.8	2.1	1.5
	Full	2.8	2.7	2.5	2.2
	Luxury	2.7	2.9	2.5	2.4
	Midsize	2.9	2.3	2.5	1.9
	Pickup	2.8	2.4	2.4	2.0
	SUV	2.7	2.6	2.5	2.2
	Sporty	2.9	2.7	2.6	2.4
	Average	3.3	3.5	3.5	3.4
Old	Compact	2.8	1.9	2.4	1.6
	Full	3.2	2.9	2.3	1.8
	Luxury	3.0	3.2	2.2	1.9
	Midsize	3.1	2.4	2.6	1.9
	Pickup	3.1	2.5	2.6	2.6
	SUV	3.0	2.8	2.5	2.0
	Sporty	3.0	2.9	2.7	2.5
	Average	3.2	3.4	3.6	3.5

Note: Quarters are defined over the automotive year, which begins in August. Cells are the averages of own-price elasticities.

estimated elasticities are slightly below those found in the previous literature; BLP, for example, report a range of elasticities between 3 and 6 at the model level, and Goldberg (1995) reports an average elasticity of 3.28. However, our approach differs from previous work in that we use transaction rather than list prices and estimate our demand system at the quarterly, rather than annual, frequency.

Given that automakers sell two vintages of the same model simultaneously for almost half of the model year, the cross-price elasticity between vintages of the same model is of particular interest. For most of the vehicles in our sample, the old and new vintages of the same model are sold simultaneously during the first and second quarters (August through January). These estimated cross-price elasticities are quite small, ranging from near 0 to 0.02 (see Appendix A for detailed numbers); various vintages of the same model, then, are typically quite imperfect substitutes.¹⁸ This result is not intuitive, given the often similar characteristics of different vintages of the same model. Yet, the implication that consumers do not consider the old and new model-year vintages as close substitutes accords with their dramatic price differences (recall the 9% new-vintage premium documented earlier). Further, from survey data, we know that households who purchase at the end of the model year typically have lower incomes compared to consumers who purchase at the beginning of the model year (see Aizcorbe, Bridgman, and Nalewaik, forthcoming). Hence, when we observe automakers selling two vintages of the same model simultaneously, there is typically a large price difference between the two vintages and the consumers purchasing each vintage have significantly different incomes. These cross-price elasticity estimates, however, are not central to our analysis; later in this article, we show our main results are robust to larger values of this cross-price elasticity.

Finally, we turn to the own-variety elasticities implied by the model. As shown in Table 6, variety plays an important role in consumers' automobile purchasing decisions. Over the first four quarters of the model's product life, increases in variety significantly bolster demand. Over this period, a 1% increase in variety bolsters sales by roughly 0.5%. The elasticities drop in the fifth and sixth quarters, however, implying there are only small gains to increasing variety at the end of the model year.

¹⁸ Ana Aizcorbe suggested that geographical factors may explain our low cross-price elasticity estimates. If different vintages of the same model are rarely offered for sale at the same location, then the degree to which consumers can substitute between vintages may be limited.

TABLE 6 Average Own-Variety Elasticities by Market Segment, Quarter, and Vintage

Vintage	Market Segment	1st Quarter	2d Quarter	3d Quarter	4th Quarter
New	Compact	0.49	0.50	0.50	0.46
	Full	0.53	0.55	0.53	0.49
	Luxury	0.52	0.51	0.53	0.54
	Midsized	0.52	0.53	0.54	0.52
	Pickup	0.52	0.53	0.54	0.52
	SUV	0.54	0.55	0.53	0.49
	Sporty	0.50	0.50	0.52	0.45
	Average	0.53	0.74	0.75	0.75
Old	Compact	0.14	0.12	0.24	0.34
	Full	0.13	0.12	0.21	0.57
	Luxury	0.12	0.12	0.19	0.19
	Midsized	0.14	0.11	0.21	0.38
	Pickup	0.13	0.10	0.16	0.55
	SUV	0.13	0.11	0.19	0.50
	Sporty	0.19	0.17	0.12	0.32
	Average	0.45	0.09	0.24	0.81

Note: Quarters are defined over the automotive year, which begins in August. Cells are the averages of own-variety elasticities.

We use these results to parameterize a reduced-form demand curve, equation (1), for each market segment. Because we are modelling the firm at a weekly frequency but have quarterly estimates, we interpolate to create elasticities at the weekly frequency. From our data, we construct a monthly time series of average price, quantity, and variety of the new and old vintage by market segment, which we interpolate to produce a weekly series. For every week, we then solve for the demand curve's constant term, μ_t^j , by assuming that the observed average price-quantity pairs for period t and market segment j , given variety and the competing vintage's price, is a point on the reduced-form demand curve. The end results are weekly demand curves for an average vehicle over its life cycle.

An important feature of the resulting sequence of static demand curves is their steady leftward shift roughly 6 months after a vehicle has been introduced. This implies that starting half a year into the product cycle, the firm faces a weakening of demand (i.e., μ_t^j is decreasing in t) over the remainder of the product cycle.

Another approach would be to use the discrete-choice demand system directly. For this alternative approach, we would use specific model data (e.g., Chrysler's Grand Jeep Cherokee) rather than our current approach of averaging across all vehicles within a market segment. We decided against this alternative approach for three reasons. First, our goal is to make a general point regarding the relationship among inventories, prices, and sales. By analyzing specific vehicles, we worried our results would not be general enough because of the idiosyncratic variation at the model level.¹⁹ Our second concern was how well the supply-side model, which does not have any supply-side shocks, could match a time series of prices, sales, and inventories of a specific vehicle. Unlike for the average car, these time series are quite variable at the vehicle level (see, for example, Figure 1). Third, we found that the log-log demand curve provided a close approximation to the discrete-choice demand curve when considering small changes in price and variety. As shown in Appendix A, there are not large differences between each approach's predicted sales for a particular model. As such, for the purposes of this article, there seems to be little cost to employing the parameterized reduced-form demand curve in place of the discrete-choice demand system.

¹⁹ For example, our data include the Ford Explorer/Firestone tire recall in 2000, the total overhaul of the Ford F-series in 2003–2004, the replacement of the popular compact Ford Escort by the compact Ford Focus, and so forth.

TABLE 7 Supply-Side Moments

Market Segment	Sales (Units)		Days-Supply		Prices					
					Average (\$)		Decline (%)		Vin. Prem. (%)	
	Data	Model	Data	Model	Data	Model	Data	Model	Data	Model
Compact	8,614	8,423	73	73	13,644	13,612	9.5	4.0	7.8	7.0
Midsize	7,760	7,178	65	65	19,063	19,023	9.2	9.1	8.1	8.6
Full-size	4,729	4,321	75	75	23,724	23,762	8.9	8.3	8.5	5.9
Luxury	2,548	2,420	80	80	35,758	35,703	12.1	11.0	12.0	8.6
Pickup	24,962	24,967	84	84	23,386	23,509	9.6	9.3	12.0	10.9
SUV	8,327	8,792	75	75	28,529	28,554	8.2	10.2	7.3	7.2
Sporty	4,239	5,234	83	83	25,887	25,919	4.9	8.7	6.4	6.5
Average	11,990	11,962	75	75	23,343	23,369	9.0	9.0	9.0	8.5

Note: Vin. Prem. stands for vintage premium, and the percentage price declines are at annual rates. Sales are the average monthly unit sales for a vehicle. Some of the data numbers presented here are also reported in earlier tables. The model numbers were computed from simulated data.

□ **Supply-side parameters.** To parameterize the cost function, we set the line speed, workers per shift, and wage rates to values typically observed at assembly plants. The line speed at most North American assembly plants is set between 35 and 60 cars per hour; thus, we fix the line speed to 45 cars per hour.²⁰ Using the employment data from Hall (2000), we set n to 1300 workers per shift, so the plant employs 2600 workers. We read the wages off the union contract: $w_1 = \$27.00$ per hour, and $w_2 = \$28.35$ per hour. Also based on the union contract and industry practices, we impose mild seasonality on production assuming that the plant closes for two weeks in July (weeks 51 and 52) for a model changeover, for a week between Christmas and New Year's Day (week 23), and for single days throughout the year corresponding to traditional holidays.

We set the remaining two parameters, γ and $1/(1+r)$, to match for each market segment two first moments in the data: the average retail price and days-supply. Although we would have preferred to formally estimate these parameters, the time needed to compute the model's solution made this infeasible. The per vehicle material cost, γ , effectively scales the cost function linearly. We set γ between 39.5% (sport) and 63% (luxury) of the average retail price to match the observed prices. We choose values of $1/(1+r)$, the weekly discount factor, between 0.962 for pickups and 0.982 for sport cars to match the average days-supply of inventories observed in the data. These values imply a high degree of impatience on the part of the automaker. A discount factor of 0.975 on a \$23,000 vehicle implies a weekly holding cost of \$575. At first blush this cost may seem high, but this parameter is the sole cost of holding inventories and thus it incorporates all the holding costs (e.g., the opportunity cost of funds, physical storage costs, insurance, physical depreciation, bookkeeping costs, etc.) that are not explicitly modelled. The parameter values for each market segment are reported in Appendix B.

5. Results

■ Given our choices of γ and r , the model closely matches the days-supply and average price across all market segments (see Table 7). The model also generates average sales that are similar to the data, although the model underpredicts midsize and full-size sales and overpredicts sporty sales. As a measure of the model's goodness of fit, we compare the model's predictions of average price decline and vintage premia against the data (the last four columns of Table 7). Overall,

²⁰ In order to match the high level of monthly sales for pickups, we set its line speed to 90. Unlike cars, which are typically assembled at only one or two plants, several popular pickup trucks (e.g., Ford F-series, Chevy Silverado, and Dodge Ram) are produced at four or five plants.

the model performs well. For four of the seven market segments (midsize, full-size, luxury, and pickups), the implied price declines are within a single standard error of the average declines seen in the data; for the sporty segment, the average decline is within the two-standard-error band.²¹ The average price decline from the model is 9.0%, matching exactly the average price decline in the data. Although the implied vintage premia for five of the seven market segments is within two standard deviations of the observed values, the model underestimates the average vintage premia slightly, 8.5% versus 9.0%. Although this is outside the two-standard-error band, we believe that relaxing our assumption that new-vintages arrive strictly every 52 weeks would enable the model to better match this moment. Overall, then, our parameterized model matches well the observed average sales, days-supply, and prices. The model’s predictions of the price decline and new-vintage premium are close to those seen in the data, demonstrating goodness of fit.

As a robustness check, we re-solved our model with a higher cross-price elasticity for different vintages of the same model. Recall our demand-side model estimates this cross-price elasticity to be essentially zero. To determine how important this parameter is to our main results, we recomputed Table 7 using a cross-price elasticity of 0.2, holding everything else constant. We chose 0.2 because this value is typically an upper bound on the demand model’s estimates of a vehicle’s cross-price elasticities.²² Reassuringly, our results are robust to the higher cross-price elasticity parameter (see Appendix B for details). Our results are also robust to larger own-price elasticities; in preliminary work, we used own-price elasticities ranging from 6 to 10 and found the same qualitative results reported in this article.²³

In the model there is a tight connection between prices and inventories. The manufacturer trades off the gain from selling a vehicle today at a specific price against the option value of having that vehicle in inventory and available for sale in the future. The firm’s optimal pricing rule, equation (14), formally lays out this tradeoff (for clarity we have set the cross-price elasticities equal to zero),

$$p_t^{this} = \frac{-s_t^{this}(p_t)}{\partial s_t^{this}(p_t)/\partial p_t} + \frac{1}{1+r} V_2(I_t^{last} - s_t^{last}, I_t^{this} + q_t^{this} - s_t^{this}, T + 1), \quad (14)$$

where V_2 , the shadow value of inventories, denotes the derivative of the value function with respect to the second argument (i.e., the inventory stock of the current vintage).

The shadow value of inventories reflects two benefits to the firm from holding inventories. An additional unit of inventory is valuable because it increases both the variety of products available to consumers and the firm’s ability to optimally schedule production. Naturally, however, these benefits are worth less when the firm already has a large stock of inventories. Figure 5 plots the shadow value of inventories in week 27 (other weeks are qualitatively similar) of at each point in the state space. As the inventory stock of this model year’s vintage increases, the shadow value of an additional unit of inventory falls. For this particular week, the range in value is from a high of \$13,827 to a low of \$11,919. Indeed, when the firm holds 50,000 of this model year’s vintage in stock, our model estimates that the marginal vehicle in inventory is worth less to the firm than the average cost of producing a vehicle running two 40 hour shifts.

This curvature in the shadow value of inventories is reflected in the automaker’s optimal pricing rule, as laid out in equation (14). We plot the pricing rule for this year’s vintage for week 26 in Figure 6 for every point in the state space. Holding the old model year’s inventory stock


²¹ Standard errors are reported in Tables 1 and 2.

²² Schiraldi (2010) reports the average cross-price elasticity between vehicles of the same type to be between 0.08 and 0.28. Schiraldi also reports cross-price elasticities between different vintages of vehicles of the same type, and these elasticities are much lower than 0.2.

²³ In fact, having higher own-price elasticities increased by how much the firm’s inventory management strategy explained the price declines observed within the model year.

FIGURE 5

WEEK 27 SHADOW VALUE OF INVENTORIES FOR THIS YEAR'S VINTAGE



Note: This figure plots the derivative of the firm's value function with respect to the current model year's inventories in week 27 of the product cycle, for many points in the state space. The state variables are the inventory stocks of the current and the old vintage.

constant, the optimal price for the new model-year vehicle substantially decreases with increases in the new model-year's stock of inventory.²⁴


This tight link between price and inventories has a substantial impact on the automaker's problem over the product cycle. To illustrate the relationship among sales, prices, inventories, and production over the product cycle, in Figure 7, we plot a simulation from the model for five 52 week model years, time aggregated to a monthly frequency. Because the model is deterministic, each of these simulations is identical. These graphs are analogous to figures presented in Section 2, however, note that Figure 12 is for a particular midsize car whereas we parameterize the model for an average midsize car. Just as we see in the data, prices decline over the model year while sales and inventories follow a hump-shaped path.

The two driving forces behind these patterns in the data are: (i) inventories' aforementioned roles in strengthening demand and allowing the firm to bunch production and (ii) the weakening of demand for a vehicle over the last two thirds of its product cycle (i.e., the constant term in equation (1), μ_t^j , decreases). Early in a vintage's product cycle, inventories are low and so the shadow value of additional units of inventories is high. Consequently, the automaker both sets a high rate of production and chooses a high price level. Because the resulting rate of sales is below the rate of production, inventories accumulate. As seen in Figure 7D, the automaker rapidly builds up inventories over the first year of a vehicle's product cycle, accumulating an enormous stock of over 30,000 vehicles 6 months after a vehicle's introduction.

²⁴ These price rules are consistent with the findings of Zettelmeyer, Scott Morton, and Silva-Risso (2006) that the average retail price at a dealership with ample inventory is about \$250 per car less relative to a dealership with low inventory.

FIGURE 6

WEEK 26 OPTIMAL PRICING RULE FOR THIS YEAR'S VINTAGE



Note: This figure plots the profit-maximizing week-26 price for the current vintage for many points in the state space. The state variables are the inventory stocks of the current and the old vintage.

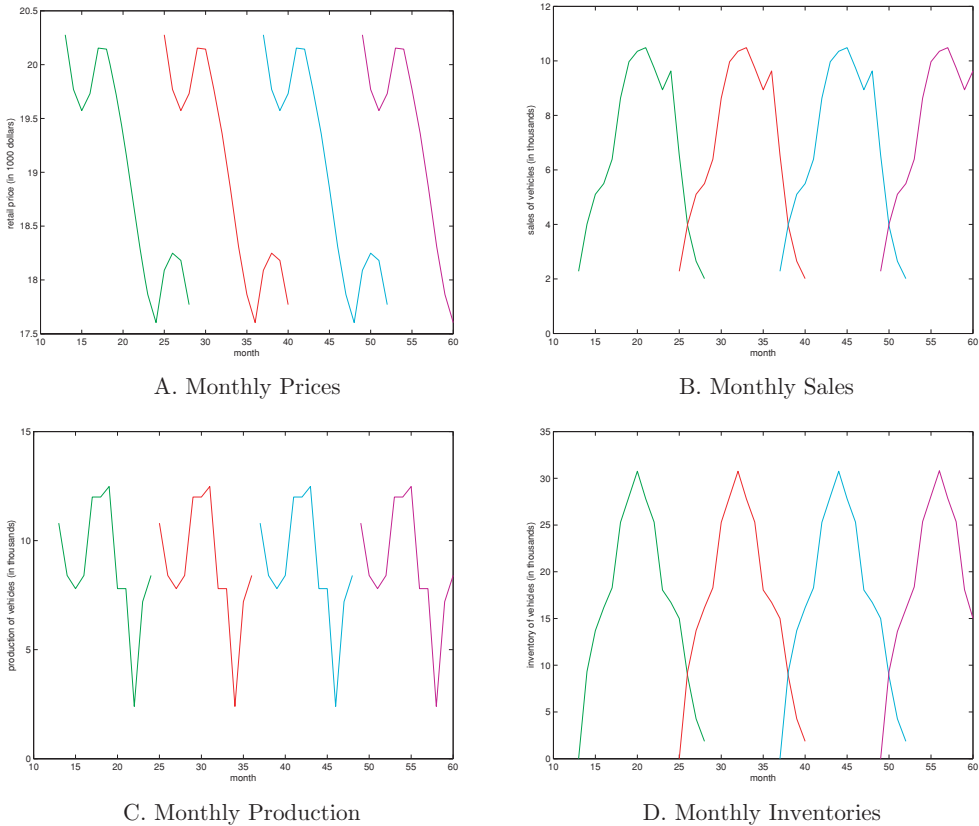
With the buildup in inventories, the shadow value of additional units of inventory falls. Further, roughly half a year into the product cycle, the firm faces a weakening of demand, which will continue throughout the product cycle. In response to this combination of forces, the firm reduces the rate of production and lowers the price. Hence, sales remain roughly constant and inventories slowly deaccumulate. The fall in demand accelerates over the last third of the product cycle, coinciding with the introduction of the next model-year's vehicle. Thus, we see both sales and prices decline over the last third of the product cycle.

As we see in the data (e.g., Figures 1A and 1C), there is considerable jaggedness in the time paths of both production and prices. The model's simulations of these two series also exhibit this pattern. The jaggedness of the simulated production series derives from the interplay of the time aggregation and the bunching of production at the weekly frequency. Due to the nonconvexities in the cost function, during periods in which the firm wishes to operate below its minimum efficient scale it minimizes costs by operating two 40 hour shifts one week and shutting down completely another week. This all-on/all-off production pattern in weekly output can generate large swings in monthly output.

The nonmonotonicity of the price contour over the product cycle reflects the impact of inventories on both demand and marginal cost. Early on in the product cycle, inventories are growing rapidly. Consequently, the demand curve is shifting out because variety is increasing, and the marginal cost curve is shifting down because the opportunity cost of selling a car (i.e., the marginal benefit of having an additional vehicle in inventory) is falling. Sales unambiguously increase, but whether prices rise or fall depends on which effect dominates. A similar process occurs at the end of the product cycle as the inventory stock dwindles, reducing demand and increasing the marginal cost of selling an additional vehicle.

FIGURE 7

SIMULATED PRICES, SALES, PRODUCTION, AND INVENTORIES FOR AN AVERAGE MIDSIZE CAR BY MODEL YEAR UNDER THE BUILD-TO-STOCK INVENTORY POLICY



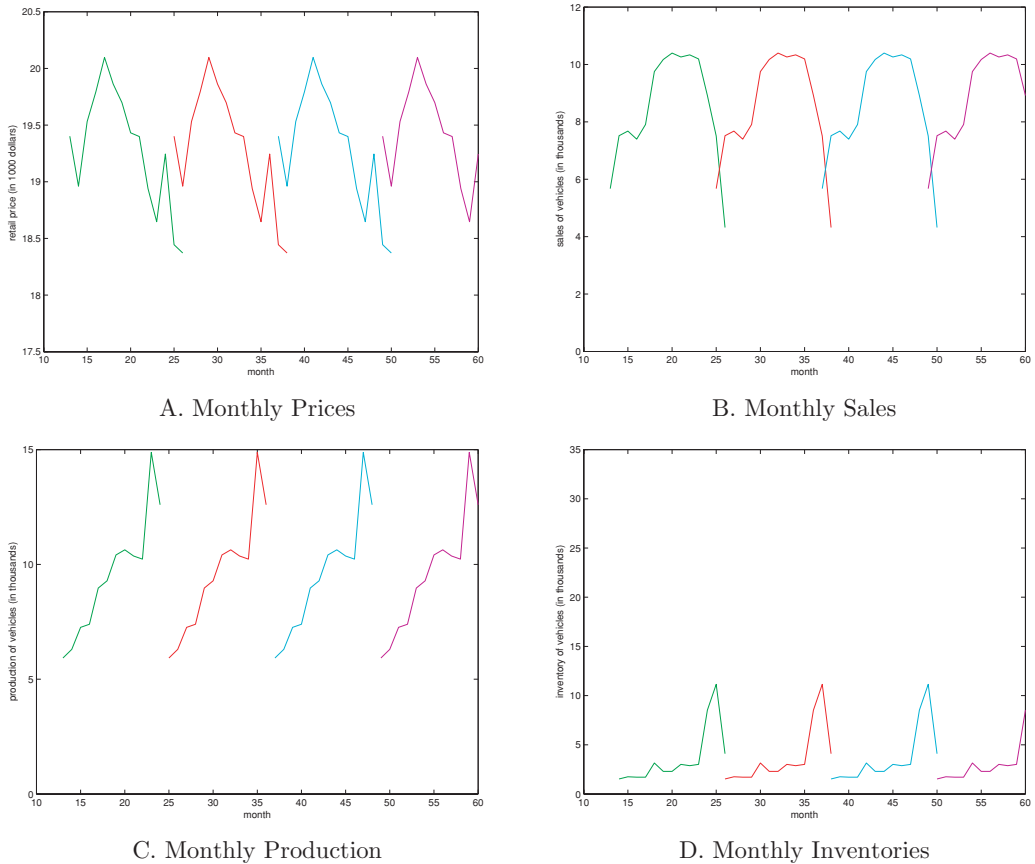
6. The counterfactual

■ To better understand the importance of inventories in stimulating demand, we employ the following counterfactual. We re-solve the model, setting the variety term in the demand curve (equation (1)) to 1.25 for all 104 weeks, turning off inventory's effect on demand. For a typical vehicle, variety peaks a little above 1.25, midway through the product cycle. Although sales still can only be made from beginning-of-period inventories (equation (3)), we interpret this simulation as approximating a “build-to-order” inventory policy in which consumers can purchase a vehicle with the exact specification they want. Hence, we assume that firms can offer consumers “full variety” throughout the product cycle without holding large levels of inventories. In this case, inventories only serve to facilitate the manufacturer's cost-minimizing production schedule and to allow the firm to sell the current vintage beyond the 12 month production period. This contrasts to the “build-to-stock” inventory management strategy firms currently use, where dealer inventories also provide variety, thereby helping match consumers to their ideal vehicle.

The counterfactual illustrates the importance of the variety-increasing role for inventories. In Figure 8, we plot monthly prices, sales, production, and inventories under the counterfactual build-to-order policy. For ease of comparison, we use the same scales as those in Figure 7, which show the model's simulation under the observed build-to-stock inventory management policy. Without the variety-increasing effect, firms hold much less inventory; the average ratio

FIGURE 8

SIMULATED PRICES, SALES, PRODUCTION, AND INVENTORIES FOR AN AVERAGE MIDSIZE CAR BY MODEL YEAR UNDER THE BUILD-TO-STOCK INVENTORY POLICY



of inventories to sales is one-fifth the level compared to the build-to-stock case (see Table 8).²⁵ In this counterfactual, the marginal benefit of inventory is roughly equal to its cost of production for the first 10 months of the product cycle. Hence, the firm sets prices such that the rate of sales closely mirrors the desired rate of production, resulting in little growth in inventories. At the end of the production cycle, the value of inventories rises because it allows the firm to sell vehicles beyond the 12 month production period. Consequently, the automaker ratchets up output and increases prices. This dampens sales and allows for a modest accumulation in inventory. Note the spike in prices, production, and inventories in Figures 8A, 8C, and 8D, respectively, 12 months into the product cycle.

In Table 8, we report differences between sales, days-supply, and prices under the built-to-stock and build-to-order inventory strategies. The contour of prices over the model year are strikingly different across these two cases. Under the counterfactual build-to-order inventory strategy, prices decline a bit more than half as much as those observed under the build-to-stock case (see also Figures 7A and 8A). The change in pricing strategy reflects the fact that the firm no longer wishes to rapidly accumulate inventories at the beginning of the product cycle. The

²⁵ As is well understood in the inventory literature, it is difficult to match the high level of inventories observed in many industries only relying on inventory's role in minimizing production costs (e.g., Bils and Kahn, 2000).

TABLE 8 Counterfactual Results

Market Segment	Sales (Units)		Days-Supply		Prices				Vin. Prem. (%)	
					Average (\$)		Decline (%)			
	BtS	BtO	BtS	BtO	BtS	BtO	BtS	BtO	BtS	BtO
Compact	8,423	9,513	73	13	13,612	13,801	4.0	1.0	7.8	1.8
Midsized	7,178	9,521	65	11	19,023	19,372	9.1	4.7	8.6	4.5
Full-size	4,321	4,997	75	14	23,762	24,017	8.3	5.9	5.9	5.6
Luxury	2,420	2,845	80	11	35,703	35,923	11.0	10.8	8.6	9.7
Pickup	24,967	25,180	84	21	23,502	24,218	9.3	5.3	10.9	3.5
SUV	8,792	9,237	75	13	28,554	28,823	10.2	6.7	7.2	6.5
Sporty	5,234	4,238	83	16	25,919	26,732	3.6	3.5	6.5	3.0
Average	11,962	12,849	75	15	23,367	23,775	9.0	5.3	8.5	4.7

Note: Vin. Prem. stands for vintage premium and the percentage price declines are at annual rates. Sales are the average monthly unit sales for a vehicle. BtS and BtO stand for build-to-stock and build-to-order, respectively. All numbers were computed from simulated data. The BtS numbers are also reported in Table 7.

counterfactual suggests then that four tenths of the overall price decline over the model year observed in the data is driven by automakers' build-to-stock inventory strategy. Further, as a result of the small price declines, the vintage premia under the build-to-order strategy are also significantly smaller relative to those under the build-to-stock case (see the last two columns of Table 8).

The time path of production is less jagged in the build-to-order case (Figure 8C) than in the build-to-stock case (Figure 7C). In both cases, the firm engages in weeklong shutdowns when desired production is less than the minimum efficient scale. However, in this particular build-to-order simulation, during the first couple of months while desired production is below the minimum efficient scale, production follows a 4–5 week cycle which is averaged out through the time aggregation. In the build-to-stock case, there is less periodicity in the shutdowns and more overtime hours are employed when the plant is operating, hence the jaggedness in the weekly data is not eliminated, but magnified, with time aggregation.

7. Conclusion

■ We have documented a set of stylized facts for the within-model-year pricing and sales of new automobiles. Prices decline steadily over the model year while sales and inventories are hump shaped. It is not the case that prices only fall during the overlap period between vintages when dealers shout over the radio, “We are slashing prices to make room for the new model year!” To understand these facts, we formulate and solve an industry model for a single vehicle line. Our model provides a consistent explanation of these facts and, through the counterfactual, highlights the role of inventories in boosting demand by increasing variety. Indeed, the model predicts this channel is important enough that it accounts for four tenths of a vehicle's price decline over the product cycle and quintuples the average inventory-to-sales ratio a firm maintains.

Advances in production and information technology have made it easier to implement build-to-order policies. For example, the computer maker Dell has been successful in selling built-to-order computers. It is our understanding from discussions with industry executives that the automakers would like to move toward an inventory policy in which a larger fraction of consumers order their new vehicles rather than buy whatever is on the dealer's lot. Our analysis suggests that enacting such a policy will dampen the within-model-year price declines and reduce the period in which consecutive vintages compete with each other.

A shortcoming of this article is the use of static demand. Although it is difficult to predict the impact of using a dynamic demand model on our results, we conjecture that our main results

would remain qualitatively the same. Hendel and Nevo (2006, 2010) show that dynamic demand models typically produce lower own-price and higher cross-price elasticities relative to static models. Furthermore, Copeland (2009) estimates a dynamic model of within-model-year new-vehicle demand and reports lower own-price elasticity estimates compared to the ones reported in this article. A change in price elasticities should not be problematic, however, because our robustness checks show that our main results hold for a range of elasticity estimates. Of course, solving the dynamic firm's problem when faced with a dynamic demand system which allows for intertemporal substitution may produce unexpected results. Incorporating dynamic demand would be particularly insightful, because we could compute how much of the observed within-model-year price declines for automobiles is due to intertemporal price discrimination, inventory management, and other forces. It is possible that our main result that automakers' inventory management accounts for four tenths of the price decline would be greatly diminished given dynamic demand. Given the results from the inventory literature, however, we find this unlikely. As mentioned in our literature review, the macroinventory literature finds that inventory's role in smoothing production is simply not enough to explain the high levels of automobiles held in inventory. Hence, in order to match the inventory levels observed in the data, this literature has relied on ad hoc assumptions that inventories play a role in generating sales. Reassuringly, this channel is supported by the industry wisdom that putting more cars on automobile dealers' lots (i.e., building up inventories) is believed to generate more sales by better matching consumers to their ideal vehicle. Consequently, we find it unlikely that a richer model incorporating dynamic demand would wipe out a role for inventories in generating sales. Given inventory's role in generating sales, our model predicts that the firm's inventory management policy will drive part of the within-model-year price declines observed in the data.

Appendix A

- In this appendix, we provide our estimated cross-price elasticities and comparisons of the sales predictions of the reduced-form demand function we use versus the sales predictions of the discrete-choice demand system.
- **Demand estimation.** Table A1 lists the estimated cross-price elasticities between vintages of the same model.
- **Comparisons of sales predictions.** Here we compare the sales predictions for specific vehicles by the log-log

TABLE A1 Cross-Price Elasticities between Vintages of the Same Model by Market Segment and Quarter

Vintage	Market Segment	1st Quarter	2d Quarter	3d Quarter	4th Quarter
New to old	Compact	0.02	0.01	0.01	0.02
	Full	0.01	0.00	0.01	0.01
	Luxury	0.02	0.01	0.00	0.01
	Midsized	0.02	0.01	0.00	0.01
	Pickup	0.04	0.01	0.00	0.17
	SUV	0.02	0.01	0.01	0.02
	Sporty	0.01	0.00	0.00	0.01
Old to new	Compact	0.01	0.03	0.02	0.01
	Full	0.01	0.02	0.01	0.01
	Luxury	0.01	0.03	0.01	0.01
	Midsized	0.01	0.02	0.01	0.01
	Pickup	0.03	0.07	0.02	0.03
	SUV	0.01	0.03	0.01	0.03
	Sporty	0.01	0.03	0.01	0.00

Note: "New to old" indicates the percentage change in the market share of the newer vintage of a model given a percentage change in the price of the older vintage. "Old to new" indicates the opposite relationship. Cells are the averages of cross-price elasticities, which are only defined when a new and old vintage of the same model are sold simultaneously. Quarters are defined over the automotive year, which begins in August.

demand curve we use in our model and by the discrete-choice demand system we estimate. For these comparisons, we parameterized the log-log demand curve with vehicle-specific elasticities implied by the discrete-choice demand system. For three different vehicles, Table A2 show the percent difference in predicted sales between the two demand models for a given (price,variety) pair, holding everything else constant, in the first quarter of the 2003 model year. For example, consider a (price,variety) pair, where both price and variety are 10% below the levels observed in the data. For the compact car, the difference in predicted sales is -2.8% . Reassuringly, Table A2 demonstrate that the log-log demand curve sales predictions are fairly close to those from the discrete-choice demand system.

TABLE A2 Percent Difference of Predicted Sales between Discrete-Choice Demand System and Parameterized Log/Log Specification

		% Δ in Price									
		-0.1	-0.05	-0.02	-0.01	0	0.01	0.02	0.05	0.1	
% Δ in variety	-0.1	-0.028	-0.011	-0.002	0.000	0.003	0.005	0.007	0.013	0.022	
	-0.05	-0.030	-0.013	-0.004	-0.002	0.001	0.003	0.005	0.011	0.020	
	-0.02	-0.031	-0.013	-0.005	-0.002	0.000	0.003	0.005	0.011	0.019	
	-0.01	-0.031	-0.013	-0.005	-0.002	0.000	0.002	0.005	0.011	0.019	
	0	-0.031	-0.014	-0.005	-0.002	0.000	0.002	0.005	0.011	0.019	
	0.01	-0.031	-0.013	-0.005	-0.002	0.000	0.002	0.005	0.011	0.019	
	0.02	-0.031	-0.013	-0.005	-0.002	0.000	0.003	0.005	0.011	0.019	
	0.05	-0.030	-0.013	-0.004	-0.002	0.001	0.003	0.005	0.011	0.020	
	0.1	-0.028	-0.011	-0.003	0.000	0.002	0.005	0.007	0.013	0.021	
		Compact Car									
		% Δ in Price									
		-0.1	-0.05	-0.02	-0.01	0	0.01	0.02	0.05	0.1	
% Δ in variety	-0.1	-0.019	-0.004	0.004	0.007	0.010	0.012	0.014	0.021	0.032	
	-0.05	-0.027	-0.011	-0.003	0.000	0.002	0.005	0.007	0.014	0.025	
	-0.02	-0.029	-0.013	-0.005	-0.002	0.000	0.003	0.005	0.012	0.023	
	-0.01	-0.029	-0.013	-0.005	-0.002	0.000	0.003	0.005	0.012	0.023	
	0	-0.029	-0.013	-0.005	-0.003	0.000	0.003	0.005	0.012	0.023	
	0.01	-0.029	-0.013	-0.005	-0.002	0.000	0.003	0.005	0.012	0.023	
	0.02	-0.029	-0.013	-0.005	-0.002	0.000	0.003	0.005	0.012	0.023	
	0.05	-0.027	-0.011	-0.003	0.000	0.002	0.005	0.007	0.014	0.025	
	0.1	-0.021	-0.005	0.003	0.006	0.008	0.011	0.013	0.020	0.031	
		Midsize Car									
		% Δ in Price									
		-0.1	-0.05	-0.02	-0.01	0	0.01	0.02	0.05	0.1	
% Δ in variety	-0.1	-0.018	-0.006	0.001	0.004	0.006	0.009	0.011	0.019	0.032	
	-0.05	-0.023	-0.011	-0.003	-0.001	0.002	0.004	0.007	0.014	0.027	
	-0.02	-0.025	-0.012	-0.005	-0.002	0.000	0.003	0.005	0.013	0.026	
	-0.01	-0.025	-0.012	-0.005	-0.002	0.000	0.003	0.005	0.013	0.026	
	0	-0.025	-0.013	-0.005	-0.003	0.000	0.002	0.005	0.013	0.025	
	0.01	-0.025	-0.012	-0.005	-0.002	0.000	0.003	0.005	0.013	0.026	
	0.02	-0.025	-0.012	-0.005	-0.002	0.000	0.003	0.005	0.013	0.026	
	0.05	-0.024	-0.011	-0.004	-0.001	0.001	0.004	0.006	0.014	0.027	
	0.1	-0.019	-0.007	0.001	0.003	0.006	0.008	0.011	0.018	0.031	
		Sports Utility Vehicle									

Note: % Δ denotes percent change. Given the observed price-variety pair in the data, both the discrete-choice and log/log demand systems predict the observed sales. The tables show the average difference in predicted sales between the two demand systems given changes to the observed price-variety pair for three different types of vehicles.

TABLE B1 Supply-Side Parameters

Market	γ	$\frac{\gamma}{\text{Mean(Price)}}$	$\frac{1}{1+r}$
Segment	(dollars)	(percent)	
Compact	6,386	47.0	0.962
Full-size	14,284	60.1	0.975
Luxury	22,509	63.0	0.979
Midsize	11,438	60.1	0.972
Pickup	13,125	52.6	0.984
SUV	16,975	59.5	0.976
Sport	10,232	39.5	0.982

Note: (γ, r) are the vehicle material costs and discount rate, respectively, of the manufacturer. For each market segment, (γ, r) were chosen to match the model to the data (see Table 7). To provide a measure of the importance of the vehicle material costs, we report the ratio of these costs to a vehicle's average price.

TABLE B2 Supply-Side Moments: Cross-Price Elasticity = 0.2

Market Segment	Sales (Units)		Days-Supply		Average (\$)		Prices Decline (%)		Vin. Prem. (%)	
	Data	Model	Data	Model	Data	Model	Data	Model	Data	Model
Compact	8,614	8,251	73	75	13,644	13,697	9.5	10.4	7.8	12.5
Midsize	7,760	7,708	65	67	19,063	18,883	9.2	2.8	8.1	2.4
Full-size	4,729	4,230	75	75	23,724	23,944	8.9	11.4	8.5	6.6
Luxury	2,548	2,327	80	82	35,758	35,290	12.1	17.0	12.0	13.7
Pickup	24,962	24,098	84	82	23,386	23,819	9.6	20.8	12.0	30.7
SUV	8,327	8,544	75	75	28,529	28,296	8.2	8.7	7.3	12.2
Sporty	4,239	4,911	83	82	25,887	25,786	4.9	7.5	6.4	11.5
Average	11,990	11,764	75	75	23,343	23,405	9.0	11.0	9.0	14.4

Note: This table is an alternative version of Table 7. The difference is that the simulated data are generated by a model with cross-price elasticities between vintages of the same model set to 0.2.

Appendix B

■ In this appendix, we provide the full set of chosen parameters on the supply side and the details on how we solved the firm's problem. Finally, we report the details of our robustness exercise where we re-solved our model using a cross-price elasticity parameter of 0.2 between vintages of the same model.

□ **Supply-side parameters.** Table B1 lists the per-vehicle material cost parameter, γ , and discount rate, $\frac{1}{1+r}$, for each market segment.

□ **Solving the firm's problem.** Because of the nonconvexities in the cost function, we solve for both the optimal level of output and the cost-minimizing production schedule through grid search. We allow weekly production, q , to take on values between 0 and 6000 in increments of 50. The grids for D_t and S_t are set from 1 to 6 and from 0 to 2, respectively, in increments of 1. The plant is closed for the week whenever $S_t = 0$. The shift length, h_t , can take on values of 7, 8, 9, or 10. So there are up to 72 feasible production schedules to evaluate for each 121 possible levels of production.

We discretize each inventory grid into 28 points from 0 to 2.25 times the mean monthly inventory stock. The distance between grid points increases with the level of inventories. Thus, the grid points are more densely spaced in the region where the value function has more curvature. For each of the 784 inventory pairs, we maximize the right-hand side of equations (10) and (11) over each sales price and level of output. Points off the two inventory grids are approximated using bilinear interpolation.

□ **Robustness tables.** In this section, we report the results from our robustness exercise. We re-solved the model assuming a new cross-price elasticity parameter of 0.2 between vintages of the same model, and keeping all other parameters the same. With this new elasticity, the model continues to match the average sales, days-supply, and price level in the data quite well (see Table B2). The model does less well, however, in matching the price declines and vintage premiums of pickups and midsize cars. Importantly, the results from our counterfactual exercise are robust to this elasticity change (see Table B3).

TABLE B3 Counterfactual Results: Cross-Price Elasticity = 0.2

Market Segment	Sales (Units)		Days-Supply		Prices					
					Average (\$)		Decline (%)		Vin. Prem. (%)	
	BtS	BtO	BtS	BtO	BtS	BtO	BtS	BtO	BtS	BtO
Compact	8,251	10,157	75	7	13,697	13,867	10.4	1.3	12.5	5.3
Midsize	7,708	8,872	67	7	18,883	19,358	2.8	5.9	2.4	2.9
Full-size	4,230	5,051	75	16	23,944	23,997	11.4	6.5	6.6	5.4
Luxury	2,327	2,845	82	7	35,290	35,923	17.0	6.1	13.7	5.9
Pickup	24,098	30,987	82	7	23,819	23,497	20.8	10.2	30.7	7.1
SUV	8,544	10,015	75	7	28,296	28,790	8.7	1.4	12.2	0.1
Sporty	4,911	4,845	82	7	25,786	26,732	7.5	2.9	11.5	3.0
Average	11,764	14,490	75	8	23,405	23,583	11.0	5.2	14.4	3.7

Note: This table is an alternative version of Table 8. The difference is that the simulated data are generated by a model with cross-price elasticities between vintages of the same model set to 0.2.

References

- ACKERBERG, D. AND RYSMAN, M. "Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects." *RAND Journal of Economics*, Vol. 36 (2005), pp. 771–788.
- AGUIRREGABIRIA, V. "The Dynamics of Markups and Inventories in Retailing Firms." *Review of Economic Studies*, Vol. 66 (1999), pp. 275–308.
- AIZCORBE, A. AND KORTUM, S. "Moore's Law and the Semiconductor Industry: A Vintage Model." *Scandinavian Journal of Economics*, Vol. 107 (2005), pp. 603–630.
- , BRIDGMAN, B. AND NALEWAIK, J. "Heterogeneous Car Buyers: A Stylized Fact." *Economic Letters*, Vol. 109(1) (2010), pp. 50–53.
- AYRES, I. AND SIEGELMAN, P. "Race and Gender Discrimination in Bargaining for a New Car." *American Economic Review*, Vol. 85 (1995), pp. 304–321.
- BERRY, S. "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics*, Vol. 25 (1994), pp. 242–261.
- , LEVINSOHN, J. AND PAKES, A. "Automobile Prices in Market Equilibrium." *Econometrica*, Vol. 63 (1995), pp. 841–890.
- , ——— AND ———. "Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy." *American Economic Review*, Vol. 89 (1999), pp. 400–430.
- BILS, M. AND KAHN, J. "What Inventory Behavior Tells Us about Business Cycles." *American Economic Review*, Vol. 90 (2000), pp. 458–481.
- BLANCHARD, O. "The Production and Inventory Behavior of the American Automobile Industry." *Journal of Political Economy*, Vol. 91 (1983), pp. 365–400.
- BRESNAHAN, T. AND RAMEY, V. "Output Fluctuations at the Plant Level." *Quarterly Journal of Economics*, Vol. 109 (1994), pp. 593–624.
- , and REISS, P. "Dealer and Manufacturer Margins." *RAND Journal of Economics*, Vol. 16 (1985), pp. 253–268.
- BUSSE, M., SILVA-RISSO, J., AND ZETTELMEYER, F. "\$1000 Cash Back: The Pass-Through of Auto Manufacturer Promotions." *American Economic Review*, Vol. 96 (2006), pp. 1253–1270.
- CACHON, G.P. AND OLIVARES, M. "Drivers of Finished Goods Inventory in the U.S. Automobile Industry." *Management Science*, Vol. 56 (2010), pp. 1586–1604.
- CHEVALIER, J. AND GOOLSBEE, A. "Are Durable Goods Consumers Forward Looking? Evidence from the College Textbook Market." *Quarterly Journal of Economics*, Vol. 125 (2009), pp. 1853–1884.
- CONLISK, J., GERSTNER, E., AND SOBEL, J. "Cyclic Pricing by a Durable Goods Monopolist." *Quarterly Journal of Economics*, Vol. 99 (1984), pp. 489–505.
- CORRADO, C., DUNN, W., AND OTOO, M. "Incentives and Prices for Motor Vehicles: What Has Been Happening in Recent Years?" In *Price and Productivity Measurement*, edited by W. E. Diewert. Victoria, BC: Trafford Press, forthcoming.
- COPELAND, A. "The Dynamics of Automobile Expenditures." Staff Report no. 394, Federal Reserve Bank of New York, 2009.
- , SHAPIRO, A. "The Impact of Competition on Technology Adoption: An Apples-to-PCs Analysis." Staff Report no. 462, Federal Reserve Bank of New York, 2010.
- ELMAGHRABY, W. AND KESKINOCAK, P. "Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices, and Future Directions." *Management Science*, Vol. 49 (2003), pp. 1287–1309.
- FEDERGRUEN, A. AND HECHING A. "Combined Pricing and Inventory Control under Uncertainty." *Operations Research*, Vol. 47 (1999), pp. 454–475.

- GOLDBERG, P. "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry." *Econometrica*, Vol. 63 (1995), pp. 891–951.
- . "Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey." *Journal of Political Economy*, Vol. 104 (1996), pp. 622–654.
- GOWRISANKARAN, G. AND RYSMAN, M. "Dynamics of Consumer Demand for New Durable Goods." Manuscript, University of Arizona, 2010.
- HALL, G. "Non-Convex Costs and Capital Utilization: A Study of Production Scheduling at Automobile Assembly Plants." *Journal of Monetary Economics*, Vol. 45 (2000), pp. 681–716.
- HAMERMESH, D. "Labor Demand and the Structure of Adjustment Costs." *American Economic Review*, Vol. 79 (1989), pp. 674–689.
- HENDEL, I. AND NEVO, A. "Measuring the Implications of Sales and Consumer Inventory Behavior." *Econometrica*, Vol. 74 (2006), pp. 1637–1673.
- AND ———. "A Simple Model of Demand Accumulation." Manuscript, Northwestern University, 2010.
- KAHN, J. "Inventories and the Volatility of Production." *American Economic Review*, Vol. 77 (1987), pp. 667–679.
- . "Why Is Production More Volatile Than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory-Holding." *Quarterly Journal of Economics*, Vol. 109 (1992), pp. 565–592.
- KARLIN, S. and CARR, R. "Prices and Optimal Inventory Policy." In K. Arrow, S. Karlin, and H. Scarf, eds., *Studies in Applied Probability and Management Science*, Stanford, Calif.: Stanford University Press, 1962.
- LANGER, A. "Demographic Preferences and Price Discrimination in New Vehicle Sales." Manuscript, University of Michigan, 2009.
- LAZEAR, E. "Retail Pricing and Clearance Sales." *American Economic Review*, Vol. 76 (1986), pp. 14–31.
- NEVO, A. "A Practitioner's Guide to Estimation of Random Coefficients Logit Models of Demand." *Journal of Economics & Management Strategy*, Vol. 9 (2000), pp. 513–548.
- . "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica*, Vol. 69 (2001), pp. 307–342.
- PASHIGIAN, P. "Demand Uncertainty and Sales: A Study of Fashion and Markdown Pricing." *American Economic Review*, Vol. 78 (1988), pp. 936–953.
- , BOWEN, B., AND GOULD, E. "Fashion, Styling, and the Within-Season Decline in Automobile Prices." *Journal of Law and Economics*, Vol. 38 (1995), pp. 281–310.
- PESENDORFER, W. "Design Innovation and Fashion Cycles." *American Economic Review*, Vol. 85 (1995), pp. 771–792.
- PETRIN, A. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy*, Vol. 110 (2002), pp. 705–729.
- REAGAN, P. "Inventory and Price Behavior." *Review of Economic Studies*, Vol. 49 (1982), pp. 137–142.
- SCHIRALDI, P. "Automobile Replacement: A Dynamic Structural Approach." Manuscript, London School of Economics, 2010.
- SCOTT MORTON, F., ZETTELMEYER, F., AND SILVA-RISSO, J. "Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities?" *Quantitative Marketing and Economics*, Vol. 1 (2003), pp. 65–92.
- STOKEY, N. "Intertemporal Price Discrimination." *Quarterly Journal of Economics*, Vol. 93 (1979), pp. 355–371.
- WHITEN, T. "Inventory Control and Price Theory." *Management Science*, Vol. 2 (1955), pp. 61–80.
- ZETTELMEYER, F., SCOTT MORTON, F., AND SILVA-RISSO, J. "Scarcity Rents in Car Retailing: Evidence from Inventory Fluctuations at Dealerships." Working Paper no. 12177, National Bureau of Economic Research, 2006.