

Inventories and the Automobile Market *

Adam Copeland[†], Wendy Dunn[‡], and George Hall[§]

December 9, 2009

Abstract

This paper 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 percent 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.

Keywords: dynamic pricing, revenue management, discrete-choice demand estimation, build-to-stock inventory policy

JEL classification: D21, D42, E22, L11, L62

*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 two referees. Finally we thank Bob Schnorbuss for helping us obtain and interpret the data from J.D. Power and Associates. George Hall gratefully acknowledges financial support from the Alfred P. Sloan Foundation. The views expressed in the paper 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.

[†]Money and Payment Studies, Federal Reserve Bank of New York, 33 Liberty St., New York, NY 10045; e-mail: adam.copeland@gmail.com ; <http://www.copeland.marginalq.com>

[‡]Division of Research and Statistics, Board of Governors of the Federal Reserve System, Mail Stop 82, 20th and C Streets, NW, Washington DC 20551; e-mail: wdunn@frb.gov

[§]Department of Economics, Brandeis University, Waltham, MA 02454; e-mail: ghall@brandeis.edu

Two common features of durable goods markets are high levels of inventories relative to sales and declining prices over the product cycle. In this paper, we jointly consider the optimal pricing and inventory management policies for automakers, the quintessential durable goods producer. Inventories play two major roles in our model. On the firm's side, inventories help manufacturer's smooth non-convex costs of production. On the consumer's side, higher levels of inventories provide more variety, thus making it easier to match consumers with their ideal vehicle. We find a tight link between inventories and prices both through inventory's production-smoothing and variety-increasing roles. Indeed, our model predicts that automakers' build-to-stock inventory management policy is responsible for four-tenths of the 9.0 percent decline (annual rate) in prices over the model year.

To explain the covariation of prices, sales, and inventories for new automobiles over the model year, we formulate an industry model. On the consumer 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; and 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. 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.

Using these estimated preferences, we compute elasticities with respect to prices and product variety, which we use to parameterize reduced-form average demand curves for each market segment and vintage over the product cycle. Taking these demand curves as given, we turn to solving the firm's dynamic inventory 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 Like many papers in this literature, we assume that the good must be sold by a fixed deadline, but we extend the theory by allowing the firm to sell two vintages simultaneously and incorporate a realistic non-convex cost structure.

Our two-sided model provides a consistent explanation of five facts about prices, sales, and inventories over the automobile product cycle: (i) the average monthly decline in retail prices (net of rebates and incentives) is 9.0 percent at an annual rate, (ii) for about half the calendar year, automakers simultaneously

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

sell two vintages of the same model, where the older vintage sells for a 9.0 percent discount, (iii) sales and inventories are humped-shaped over the product cycle, (iv) the mean ratio of inventories to sales is 75 days, and (v) all other things being constant, higher inventories are associated with lower retail prices.

We are able to replicate facts (i) and (iii) by modeling the firm as solving an inventory control problem while facing declining demand over the product cycle. Early in the model year, the automaker sets price sufficiently high to keep sales less than production and thus accumulates 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 percent decline in the price over the entire product cycle and an average vintage premium of 8.5 percent. Because inventories are used to both smooth production and increase variety, the model is able to match the high level of inventories relative to sales (fact (iv)). Finally, because holding inventories is costly, we also find that higher inventories are associated with lower prices, all else equal (fact (v)).

To quantify the importance of inventories for the firm's decision problem, we simulate the model under a counterfactual build-to-order strategy, where firms are able to offer consumers full variety for every product without holding inventories. Under this alternative policy, we find that automakers' pricing strategies are significantly different: Within model-year prices decline by 5.3 percent, roughly six-tenths of the percent price decline observed under a build-to-stock policy. Our model therefore predicts that a firm's inventory strategy plays a central role in the firm's optimal pricing path.

Price declines over the product cycle and the simultaneous sale of multiple vintages of the same product are not unique to the automobile industry, having been documented for a number of other products, including textbooks (Chevalier and Goolsbee, 2007), microprocessors (Aizcorbe and Kortum, 2005), and consumer electronics (Gowrisankaran and Rysman, 2007; and Copeland and Shapiro, 2009). Most existing theories that seek to explain this price path focus on intertemporal price discrimination (e.g. Stokey, 1979; and Conlisk, Gerstner and Sobel, 1984) or fashion (e.g. Lazear, 1986; Pashigian, 1988; and Pendorfer, 1995). Our paper builds on this body of work by emphasizing the significant role that firm-held inventories can play in explaining the observed time-series of prices and sales.

1 Data Sources and Empirical Observations

In this section we outline our data sources and document five stylized facts.

1.1 Data Sources

To construct a dataset of transaction prices, sales, production, and inventories by model and model year in the U.S. 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 dataset was constructed by Corrado, Dunn, and Otoo (2004), using data from J.D. Power and Associates (JDPA). JDPA collects daily transaction-level information from dealerships across the U.S. 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 percent of the geographical markets in the U.S. and roughly 15 to 20 percent 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.² 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.³

We linked the JDPA dataset to a dataset 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

²A trade-in vehicle's benchmark value is the wholesale price.

³Our 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 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 BEA's personal consumption deflator.

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 dataset 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 U.S. from 1999 to 2003.

1.2 Empirical Observations

As described in the introduction, we observe several stylized facts that hold across models and model years. To provide illustrative examples, we plot in figures 1-8 the price, sales, production, and inventory data for a midsize car and a pickup truck. The steady decrease in price for each vintage is immediately evident for both vehicles, as shown in figures 1 and 5. In the 2000 model year, the average price for the midsize car falls over \$2,000, more than 10 percent of the initial price. The declines in prices for subsequent model years are just as pronounced. For the pickup, the price declines average a dramatic \$4,000 for the 2001 through 2003 model years. Both the midsize car and the pickup clearly 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 and pickup is 7 and 9 percent respectively. Turning to figures 2, 4, 6 and 8, we find the sales and inventories of both vehicles exhibit hump-shaped profiles.

Table 1 provides a summary of the average monthly price decline, the first stylized fact, by market segment and model year. For the midsize market segment, the mean monthly price decline of 1999 model-year vehicles is 9.1 percent at an annual rate. On average, midsize automobiles fall 9.2 percent. Table 1 illustrates the wide range in average price declines both across market segments and model years.⁴ 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 percent at an annual rate.

⁴We exclude the Van market segment from our analysis because a substantial number of vans are sold to firms.




Figure 1: Average Transaction Prices.




Figure 2: Monthly Sales.

The dashed line is the sum of sales across model years.




Figure 3: Monthly Production.




Figure 4: Monthly Inventories.

Prices, Sales, Production, and Inventories for a Midsize Car by Model Year

Source: J.D. Power and Associates, Ward's Communications and authors' calculations




Figure 5: Average Transaction Prices.




Figure 6: Monthly Sales.
The dashed line is the sum of sales across model years.




Figure 7: Monthly Production.




Figure 8: Monthly Inventories.

Prices, Sales, Production, and Inventories for a Pickup Truck by Model Year
Source: J.D. Power and Associates, Ward's Communications and authors' calculations




Figure 9: Matched-Model Price Indexes by 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)
Fullsize	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 parenthesis

Table 1: The Monthly Price Decline (annual rate) 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)
Fullsize	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 parenthesis

Table 2: The Average ‘New Vintage Premium’ by Market Segment and Model Year

To observe the within-year price declines more generally, figure 9 illustrates the aggregate matched-model price indexes for successive model years constructed by Corrado, Dunn, and Otoo (2004) using the entire JDPA dataset. 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. The overlap of the model-year price indexes highlights the second stylized fact: multiple vintages of vehicles are simultaneously sold for a significant portion of the model year. In our sample, the typical vehicle is on the market for 16.7 months. The number of months sold varies little across vehicles; the mean length of the automobile product cycle has a standard error of only 0.02.

Turning to the third stylized fact, we report in table 2 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 percent. Across model years, the average new vintage premium is typically between 5 and 9 percent, though the premium during the 2003 to 2004 changeover is 14.0 percent. 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 percent new vintage premium recorded for 2004 model-year pickup trucks reflects, in part, a quality improvement made to the 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 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

Variable	Coefficient	Standard Error
Age	-0.093	0.004
Odometer (thousands of miles)	-0.004	0.000

Table 3: Coefficients on Age and Odometer from the Used-vehicle Price Regression

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 differ 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 over the model year 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. We claim that 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 dataset of *used*-vehicle transactions from 2001-2003. The left-hand-side variable is the log of the transaction price for a given model and vintage of a used vehicle. The explanatory variables include time and model dummies as well as vehicle characteristics such as engine size. As a proxy for the vehicle's physical depreciation, we include the vehicle's odometer reading when sold. Finally, we also add a measure of the vehicle's model age, which equals the calendar year minus the model year plus one.⁵ As expected, the coefficient on the odometer reading is negative and implies a price decline of about 0.4 percent for each additional 1,000 miles driven (see table 3). Notably, the coefficient on age implies that, even after controlling for the odometer reading and other vehicle characteristics, increasing model age by one year decreases the value of a used vehicle by 9.3 percent, a figure only slightly greater than our estimate of the new vintage premium. This strongly suggests that the new vintage premium is driven by the difference in the new vehicles' values in the used-vehicle market.

The fourth stylized fact states that both sales and inventories for a particular model and model year exhibit a humped-shaped pattern. We measure the product cycle trend using regression analysis, by sepa-

⁵Because we have a limited set of physical characteristics to control for changes in vehicle quality across vintages of the same model, we restrict the sample to vehicles of age four or under. This restriction reduces the variation in price across vintages of the same model due to changes in unobserved characteristics.

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

Note: Standard errors are in parenthesis

Table 4: The Average Days-Supply by Market Segment

rately regressing the log of sales and log of inventories on dummy variables for each month in the product cycle, while controlling for the months of the calendar year and model fixed effects. We find the product cycle trend for both sales and inventories starts low and increases for several months before reaching a plateau. Then after almost a year into the product cycle, both sales and inventories rapidly fall.

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, or enough inventories to sell vehicles for over *three* months without any additional production! Table 4 provides a breakdown of the average days-supply by market segment and illustrates the substantial variation in days-supply across different types of vehicles.

The last stylized fact focuses on the correlation between inventories and prices. To analyze this relationship we use the residuals from the regression described above that estimated the contour of inventories over the product cycle. We compute whether prices are correlated with this measure of inventory's deviation-from-trend by regressing the log of price on a lag of the inventory residuals. We find the expected significant negative relationship between lagged inventory residuals and price: a one percent deviation in inventories above its trend is correlated with an eight percent reduction in price.

For the remainder of the paper, 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 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.

2 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 modeled 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.⁶ While issues of vertical control and price discrimination are present in the automobile market, we are implicitly assuming that manufactures and dealers jointly set prices to maximize their combined profits and solve the double-marginalization problem. 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.

The automaker sells two products: this year's vintage and last year's vintage. The decision period is a week. There are 52 weeks in a model year, and a new model year begins the week after the old model year ends. So the automaker solves an infinite horizon problem by repeatedly solving a 52-week problem. Successive model years are linked because this year's vintage becomes last year's vintage at the end of the 52nd week. Each week the firm must decide (1) the number of vehicles of the current model year to produce, q_t ; (2) the number of days to operate the plant, D_t , the number of shifts to run, S_t , and the number of hours per shift, h_t ; (3) the retail price of the current vintage, p_t^{his} ; and (4) 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,

⁶For example, Bresnahan and Reiss (1985) model and estimate the division of markups between automobile manufacturers and dealers. Busse, Silva-Risso and Zettlemeyer (2006) estimate how the value of manufacturer's incentives programs are split between dealers and final customers. For discussions of bargaining and price discrimination in the retail auto market see Ayres and Siegelman (1995), Goldberg (1996), Scott Morton, Zettlemeyer, and Silva-Risso (2003), and Langer (2009).

$\frac{I_t^j}{j_{mean}}$, 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 U.S. 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 U.S. in 2002, while Honda had 959.⁷ 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.⁸

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 (2008).

Since 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:

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

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:

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

⁷Data taken from Ward's 2002 *Automotive Yearbook*.

⁸A 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).

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

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

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:

$$I_1^{last} = I_{52}^{this} + q_{52} - s_{52}^{this}. \quad (6)$$

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 is typical in most manufacturing industries, assembly plant managers increase or decrease production by altering the workweek rather than the rate of production. 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:

$$q_t = D_t \times S_t \times h_t \times LS. \quad (7)$$

Although the production function is linear, the firm faces several important non-convexities in costs. Past work by Bresnahan and Ramey (1994), Hall (2000), and Ramey and Vine (2006), has shown the importance of accounting for the non-convex cost structure faced by automakers to explain the volatility in automobile production due to frequent temporary plant shutdowns. Hence, we explicitly model the costs firms face in producing automobiles, where these costs are a function of both the underlying technology and the negotiated labor contracts.

From the autoworkers' union contracts, we know that workers on the second shift receive a 5 percent premium above the first shift wage. Any work in excess of eight 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 percent 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

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

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 percent 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 a several weeks and then shutting down the plant for a week.

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

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

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:

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

$$+ \left(\frac{1}{1+r} \right)^{52} V(I_{52}^{this} + q_{52} - s_{52}^{this}, 0, 1) \Big\}.$$

For a given parameter vector, we carried out the following steps to solve for the fixed point: (1) Guess an initial value for $V(I^{last}, 0, 1)$; (2) solve the 52 Bellman equations in (10) and (11) through backward recursion; (3) compute a new value for $V(I^{last}, 0, 1)$ through policy iteration; and (4) repeat steps 2 and 3 until a fixed point is reached. More details on the solution method are provided in the appendix.

3 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.

3.1 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 1, our sample includes data on the Big Three firms over the five-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.⁹ 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 previously been incorporated into the BLP framework. We use the definition of variety described above (see section 2), 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.

⁹Information 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 U.S. 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.¹⁰ 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.¹¹ 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.¹² We follow BLP and use competing products' characteristics as instruments. Because characteristics vary at the model-year frequency while price and sales vary quarterly, these instruments have diminished effect. We augment these instruments with the congestion variable, which is an indicator variable equal to one when multiple vintages of a model are sold simultaneously. This variable is strongly correlated with price and varies at a quarterly frequency. Further, there is little reason to suspect that the sale of multiple vintages of a model is correlated with the unobserved characteristic of a vehicle. This is because new vintages are most often introduced at an annual frequency, and automakers face large costs to altering the scheduled introduction of a new vintage.

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} . Since 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.

Results: We present a subset of the parameter estimates in table 5. 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

¹⁰We tried an alternative approach that links the number of households per quarter to total light motor vehicle sales. We defined the market size in the first period of the model as one-fourth of households in 1999. We then used the percentage change in total light motor vehicle sales in the U.S. to grow out market size. Unlike before, with this alternative approach there is no upward trend in market size over the sample period. Further there is a correlation between the share of the outside good and the number of households looking to purchase a new vehicle. The estimated parameters and implied elasticities, however, did not significantly change with this alternative definition of market size.

¹¹We 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.

¹²See Berry (1994) for a detailed explanation of, and solution to, this problem for discrete-choice demand models.

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

Table 5: Parameter 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 5 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 Akerberg 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 drop off in its value in the fourth quarter. The quarters differ from calendar quarters. We defined the first quarter as the first three 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.

Vintage	Market Segment	1st Quarter	2nd Quarter	3rd 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

Table 6: The Average Absolute Value of Own-Price Elasticities by Market Segment, Quarter, and Vintage

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 6). 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 percent price increase for a typical compact car (roughly \$140) causes a 2.6 percent fall in sales, holding everything else equal. Own-price elasticities vary little across quarters. In general, our 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). The estimated cross-price elasticities are

Vintage	Market Segment	1st Quarter	2nd Quarter	3rd 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
	Midsize	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
	Midsize	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

Table 7: Average Own-Variety Elasticities by Market Segment, Quarter, and Vintage

quite small, ranging from near 0 to 0.02 (see the appendix for detailed numbers); various vintages of the same model, then, are typically quite imperfect substitutes.¹³ 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 percent new vintage premium documented earlier). A possible explanation for this pricing pattern is price discrimination. By setting prices such that they decline over the model year, automakers may be separating out eager consumers who are impatient for the latest and greatest vehicle, from patient consumers who are willing to wait to purchase a new vehicle at the end of the model year.¹⁴ In such an scenario, small price changes would induce little substitution across vintages of the same model.

Finally, we turn to the own-variety elasticities implied by the model. As shown in table 7, variety plays an important role in consumers' automobile purchasing decisions. Over the first 4 quarters of the model's product life, increases in variety significantly bolster demand. Over this period, a 1 percent increase in variety bolsters sales by roughly 0.5 percent. The elasticities drop in the fifth and sixth quarter, however,

¹³Ana 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.

¹⁴Supporting the claim of price discrimination in the new vehicle market, Aizcorbe et al (2007) reports survey data that shows that incomes of new vehicle purchasers falls by over 8 percent over the automotive model year.

implying there are only small gains to increasing variety at the end of the model year.

We use these results to parameterize a reduced-form demand curve, equation 1, for each market segment. Because we are modeling 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 result 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 six 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.

To provide a better sense of inventory's role in the demand curve, in figure 10 we illustrate a typical Midsize car's sales given various levels of inventories, holding everything else constant. On the right axis, we plot the linear relationship between inventories and variety. As shown in the figure, the demand curve incorporates the intuitive result that a unit in inventory strengthens sales to a larger extent when the stock of inventories is low. In addition, we plot the relationship between sales and price for a typical Midsize car in figure 11, both when it is the new vintage (week 26) and when it is the old vintage (week 60). Notice that in week 60, predicted sales are zero once price is above \$20,100, the price at which we have assumed, for this graph, the new vintage is selling in week 60 (see equation 2 for details on this demand constraint for the old vintage).

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.¹⁵ 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,

¹⁵For example, our data includes the Ford Explorer/Firestone tire recall in 2000, the total overhaul of the Ford F-series in 2003-4, the replacement of the popular compact Ford Escort by the compact Ford Focus, etc.




Figure 10: Sales, Variety and Inventory Relationship (week 26)

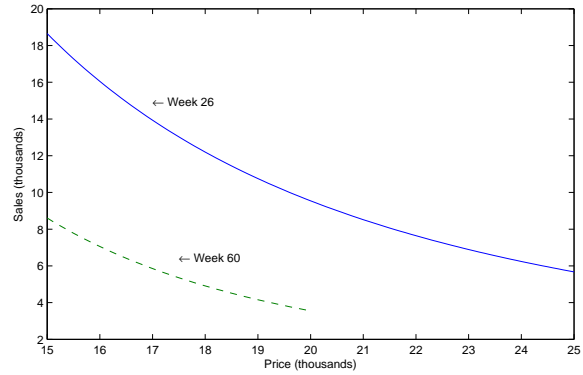


Figure 11: Sales and Price Relationship

Reduced-Form Demand Curve for the Typical Midsize Car

sales, and inventories of a specific vehicle. Unlike for the average car, these time-series are quite variable at the vehicle level. 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 the appendix, there are not large differences between each approach’s predicted sales for a particular model. As such, for the purposes of this paper, there seems to be little cost to employing the parameterized reduced-form demand curve in place of the discrete-choice demand system.

3.2 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.¹⁶ 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 Years 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

¹⁶In 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.

Market Segment	Sales		Days Supply		Prices					
	<i>Data</i>	<i>Model</i>	<i>Data</i>	<i>Model</i>	Average		Decline		Vin. Prem.	
					<i>Data</i>	<i>Model</i>	<i>Data</i>	<i>Model</i>	<i>Data</i>	<i>Model</i>
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
Fullsize	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.

Table 8: Supply-Side Moments

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 percent (sport) and 63 percent (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 to 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 of part of the automaker. A discount factor of 0.975 on \$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, book-keeping costs ...) that are not explicitly modeled. The parameter values for each market segment are reported in the appendix.

4 Results

Given our choices of γ and r , the model closely matches the days-supply and average price across all market segments (see table 8). The model also generates average sales that are similar to the data, although the model under-predicts midsize and fullsize sales and over-predicts 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 4 columns on table 8). Overall, the model performs well. For four of the seven market segments (midsize, fullsize, 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.¹⁷ 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%. While 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.

To illustrate the importance of inventories in the model, we consider the firm’s pricing decision for a typical midsize car. Because the automaker faces a downward-sloping demand curve, the profit-maximizing price sets marginal revenue equal to marginal cost. If we set the cross-price elasticities equal to zero, the optimal price for this year’s model is

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

where V_2 denotes the derivative of the value function with respect to the second argument. This is the standard condition for monopoly pricing, but in this case marginal cost is the shadow value of an additional unit of inventory next period. The opportunity cost of selling a vehicle this week is the inability to sell it next week. Hence, the shadow values of inventories drives, in large part, the firm’s optimal pricing rule.

In figure 12, we plot the shadow value of inventories for week 27 (other weeks are qualitatively similar), at each point in the state space. The shadow value of inventories is a decreasing function of the stock of inventories, ranging from \$13,827 to \$11,919. An additional unit of inventory is valuable to the firm because it increases both the firm’s ability to optimally schedule production and the variety of products available to consumers. Naturally, however, these benefits are worth less when the firm already has a large stock of inventories; when the firm holds 50,000 vehicles 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.

We then plot the pricing rule for this year’s vintage for week 26 in figure 13 for every point in the state space. As anticipated by equation (14), the pricing rule is almost the shape of the shadow value of inventories. Holding the old model year’s inventory stock constant, the optimal price for the new model year vehicle is a decreasing function of its level of inventory—our fifth fact.¹⁸

¹⁷Standard errors are reported in tables 1 and 2.

¹⁸These price rules are consistent with the findings of Zettelmeyer, Scott Morton, and Silva-Risso (2003) that the average retail




Figure 12: Week 27 Shadow Value of Inventories for This Year's Vintage.




Figure 13: Week 26 Optimal Pricing Rule for This Year's Vintage

To further illustrate the relationship among sales, prices, inventories and production, in figures 14-17, 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 1; however note that figures 1-4 are for a particular midsize car while 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 (1) inventories' role in strengthening demand and allowing the firm to bunch production and (2) the weakening of demand for a vehicle over the last two-thirds of its product cycle. Because inventories provide the automaker with the dual benefit of allowing the firm to bunch production and to increase the variety of vehicles available to consumers, firms are eager to quickly build up their stock of inventories. Indeed, the automaker sets prices high early on in the model year so that sales are below production and inventories are allowed to accumulate. As seen in figure 17, firms rapidly build up inventories over the first year of a vehicle's product cycle, accumulating an enormous stock of over 30,000 vehicles six months after vehicle's introduction.

After the buildup in inventories over the first few months of the product cycle, prices steadily decline over the middle part of the product cycle while sales remain roughly constant. These movements in prices and sales reflect the second driving force in the model, the weakening of demand over the last two-thirds of the product cycle (i.e. the constant term in equation 1, μ_t^j , decreases). The fall in demand accelerates over the end of the product cycle, coinciding with the introduction the next model year vehicle. As a price at a dealership with ample inventory is about \$230 per car less relative to a dealership with low inventory.

consequence of this rapid fall in demand, sales rapidly decline.

As we see in the data (figures 1, 3, 5, and 7), 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 comes from the bunching of production induced by the nonconvexity of the plant's cost function and plant closures for holidays. When the firm is operating below its minimum efficient scale, it minimizes costs by operating two 40-hour shifts one week and shutting down completely the following week. In the auto industry lingo such closures are called inventory adjustment shutdowns. This all-on/all-off production pattern at the weekly frequency induces non-smooth monthly time series.

The non-monotonicity 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 is falling. Sales unambiguously increase but whether prices rise or fall depends in which effect dominates. Hence the non-monotonicities in the price simulations. 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.

5 The Counterfactual

To better understand the importance of inventories in stimulating demand, we employ the following counterfactual. We resolve 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 allowing the firm to engage in 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 play the roles of facilitating the manufacturer's cost minimizing production schedule and allowing the firm to sell the current vintage beyond the twelve month production period. This stands in contrast to the "build-to-stock" inventory management strategy firms currently use, where dealer inventories also provide variety and so help match consumers to their ideal vehicle.




Figure 14: Monthly Prices.




Figure 15: Monthly Sales.




Figure 16: Monthly Production.




Figure 17: Monthly Inventories.

Simulated Prices, Sales, Production, and Inventories for a Typical Midsize Car By Model Year Under the Build-to-Stock Inventory Policy.

Market Segment	Sales		Days Supply		Prices					
	BtS	BtO	BtS	BtO	Average		Decline		Vin. Prem.	
					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
Midsize	7,178	9,521	65	11	19,023	19,372	9.1	4.7	8.6	4.5
Fullsize	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, the percentage price declines are at annual rates, BtS and BtO stand for build-to-stock and build-to-order respectively

Table 9: Counterfactual Results

The counterfactual illustrates the importance of the variety-increasing role for inventories. Without this effect, firms hold much less inventory; the average ratio of inventories to sales is one-fifth the level compared to the build-to-stock case (see table 9).¹⁹ Inventories are now mainly used to lengthen the sales period of a car. Early in the product cycle the rate of production roughly equals sales. The firm ratchets up output in the last months before production switches over to the new model year vintage. At the same time, the firm temporarily raises prices, damping sales, and so inducing a spike in inventories (see figure 21). Figures 14 and 18 illustrate that prices decline a bit more than half as much as those observed under the build-to-stock case. The change in pricing strategy reflects the fact the firm no longer wishes to rapidly accumulate inventories at the beginning of the product cycle. This simulation suggests then, that four-tenths of the overall price decline over the model year observed in 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.

6 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

¹⁹As 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 production smoothing (e.g. Bills and Kahn, 2000).




Figure 18: Monthly Prices.




Figure 19: Monthly Sales.




Figure 20: Monthly Production.




Figure 21: Monthly Inventories.

Simulated Prices, Sales, Production, and Inventories for a Typical Midsize Car By Model Year Under the Build-to-Order Inventory Policy.

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 damp the within-model-year price declines and reduce the period in which consecutive vintages compete with each other.

References

- [1] Aizcorbe, A., Kortum, S. 2005. Moore's Law and the semiconductor industry: A vintage model. *Scandinavian Journal of Economics* 107, 603-630.
- [2] Aizcorbe, A., Bridgman, B., Nalewaik J. 2007. The implications of heterogeneous buyers for measuring quality change. manuscript, Bureau of Economic Analysis.
- [3] Ayres, I., Siegelman P., 1995. Race and gender discrimination in bargaining for a new car. *American Economic Review* 85, 304-321.
- [4] Akerberg, D., Rysman, M. 2005. Unobserved product differentiation in discrete-choice models: estimating price elasticities and welfare effects. *Rand Journal of Economics* 36, 771-788.
- [5] Berry, S., 1994. Estimating discrete-choice models of product differentiation. *Rand Journal of Economics* 25, 242-261.
- [6] Berry, S., Levinsohn J., Pakes A., 1995. Automobile prices in market equilibrium. *Econometrica* 63, 841-890.
- [7] Berry, S., Levinsohn J., Pakes A., 1999. Voluntary export restraints on automobiles: Evaluating a trade policy. *American Economic Review* 89, 400-430.
- [8] Bils, M., Kahn J. 2000. What inventory behavior tells us about business cycles. *American Economic Review* 90, 458-481.
- [9] Bresnahan, T., Reiss, P., 1985. Dealer and manufacturer margins. *Rand Journal of Economics* 16, 253-268.
- [10] Bresnahan, T., Ramey, V., 1994. Output fluctuations at the plant level. *Quarterly Journal of Economics* 109, 593-624.
- [11] Busse, M., Silva-Risso, J., Zettlemeyer, F., 2006. \$1000 cash back: The pass-through of auto manufacturer promotions. *American Economic Review* 96, 1253-1270.
- [12] Cachon, G.P., Olivares, M., 2008. Drivers of finished goods inventory in the U.S. automobile industry, SSRN working paper.

- [13] Chevalier, J., Goolsbee A., 2007. Are durable good consumers forward looking? Evidence from the college textbook market, manuscript, Yale University.
- [14] Conlisk, J., Gerstner, E., Sobel. J., 1984. Cyclic pricing by a durable goods monopolist. *Quarterly Journal of Economics* 99, 489-505.
- [15] Corrado, C., Dunn, W., Otoo M., 2004. An initial look at incentives and prices for motor vehicles: What has been happening in recent years? manuscript, Board of Governors of the Federal Reserve.
- [16] Copeland, A., Shapiro, A., 2009 The impact of competition on technology adoption: An Apples-to-PCs analysis. manuscript, Federal Reserve Bank of New York.
- [17] Elmaghraby W., Keskinocak P. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science* 49, 1287-1309.
- [18] Federgruen, A., Heching A., 1999. Combined pricing and inventory control under uncertainty. *Operations Research* 47, 454-475.
- [19] Goldberg, P., 1995. Product differentiation and oligopoly in international markets: The case of the U.S. automobile industry. *Econometrica* 63, 891-951.
- [20] Goldberg, P., 1996. Dealer price discrimination in new car purchases: Evidence from the consumer expenditure survey. *Journal of Political Economy* 104, 622-654.
- [21] Gowrisankaran, G., Rysman, M., 2007. Dynamics of consumer demand for new durable goods. manuscript, University of Arizona.
- [22] Hall G., 2000. Non-convex costs and capital utilization: A study of production scheduling at automobile assembly plants. *Journal of Monetary Economics* 45, 681-716.
- [23] Kahn J., 1987. Inventories and the volatility of production. *American Economic Review* 77(4), 667-679.
- [24] Kahn J., 1992. Why is production more volatile than sales? Theory and evidence on the stockout-avoidance motive for inventory-holding. *Quarterly Journal of Economics* 109(3), 565-592.
- [25] Karlin, S., Carr R., 1962. Prices and optimal inventory policy. in *Studies in Applied Probability and Management Science*, K. Arrow, S. Karlin, H. Scarf (ed.), Stanford, CA: Stanford University Press.

- [26] Langer, A., 2009. Demographic preferences and price discrimination in new vehicle sales. manuscript, University of California.
- [27] Lazear, E., 1986. Retail pricing and clearance sales. *American Economic Review* 76, 14-31.
- [28] Nevo, A., 2000. A practitioner's guide to estimation of random coefficients logit models of demand. *Journal of Economics & Management Strategy* 9, 513-548.
- [29] Nevo, A., 2001. Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2), 307-342.
- [30] Pashigian, P., 1988. Demand uncertainty and sales: A study of fashion and markdown pricing. *American Economic Review* 78, 936-953.
- [31] Pashigian, P., Bowen B., Gould E., 1995. Fashion, styling, and the within season decline in automobile prices. *Journal of Law and Economics* 38, 281-310.
- [32] Pesendorfer, W., 1995. Design innovation and fashion cycles. *American Economic Review* 85, 771-792.
- [33] Petrin, A., 2002. Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy* 110, 705-729.
- [34] Ramey, V., Vine D., 2006. Declining volatility in the U.S. automobile industry. *American Economic Review* 96, 1876-1889.
- [35] Scott Morton F., Zettelmeyer, F., Silva-Risso J., 2003. Consumer information and discrimination: Does the internet affect the pricing of new cars to women and minorities? *Quantitative marketing and Economics* 1, 65-92.
- [36] Stokey, N., 1979. Intertemporal price discrimination. *Quarterly Journal of Economics* 93, 355-371.
- [37] Whiten, T., 1955. Inventory control and price theory. *Management Science* 2, 61-80.
- [38] Zettelmeyer, F., Scott Morton F., Silva-Risso J., 2003. Inventory fluctuations and price discrimination: The determinants of price variation in car retailing. manuscript, Yale School of Management.

7 Appendix

In this appendix, we provide: (1) our estimated cross-price elasticities, (2) comparisons of the sales predictions of the reduced-form demand function we use versus the sales predictions of the discrete-choice demand system,(3) the full set of chosen parameters on the supply side, and (4) details on how we solved the firm’s problem.

7.1 Demand estimation

Table 10 lists the the cross-price elasticities between vintages of the same model.

Vintage	Market Segment	1st Quarter	2nd Quarter	3rd 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
	Midsize	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
	Midsize	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

Notes: “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.

Table 10: Cross-Price Elasticities Between Vintages of the Same Model by Market Segment and Quarter

7.2 Comparisons of Sales Predictions

Here, we compare the sales predictions for specific vehicles by the log-log 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. Each table shows 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 percent below the levels

observed in the data. For the compact car, the difference in predicted sales is -2.8 percent. Reassuringly, the tables below demonstrate that the log-log demand curve sales predictions are fairly close to those from the discrete-choice demand system.

7.3 Supply-side parameters

Table 14 lists the per-vehicle material cost parameter, γ , and discount rate $\frac{1}{1+r}$, for each market segment.

7.4 Solving the firm's problem

Because of the non-convexities 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 bi-linear interpolation.

		% Δ in price								
		-0.1	-0.05	-0.02	-0.01	0	0.01	0.02	0.05	0.1
% Δ in vari- ety	-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

Table 11: Compact car

		% Δ in price								
		-0.1	-0.05	-0.02	-0.01	0	0.01	0.02	0.05	0.1
% Δ in vari- ety	-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

Table 12: Midsize car

		% Δ in price								
		-0.1	-0.05	-0.02	-0.01	0	0.01	0.02	0.05	0.1
% Δ in vari- ety	-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

Table 13: Sports utility vehicle

Percent Difference of Predicted Sales Between Discrete-choice Demand System and Parameterized Log/log Specification

Market Segment	γ (dollars)	$\frac{\gamma}{\text{mean}(\text{price})}$ (percent)	$\frac{1}{1+r}$
Compact	6,386	47.0	0.962
Full	14,284	60.1	0.975
Luxury	22,509	63.0	0.979
Midsized	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

Table 14: Supply-side Parameters