



Yale University
Department of Economics

Yale Working Papers on Economic Applications and Policy

Yale University
P.O. Box 208268
New Haven, CT 06520-8268
DISCUSSION PAPER NO. 2

**Prices, Production and Inventories over
the Automotive Model Year***

Adam Copeland
Board of Governors of the Federal Reserve System

Wendy Dunn
Board of Governors of the Federal Reserve System

George J. Hall
Yale University
Department of Economics

March 2005

*We want to thank Ana Aizcorbe, Steve Berry, Andrew Cohen, Gautam Gowrisankaran, Amil Petrin, John Rust, John Stevens, and seminar participants at numerous conferences and seminars for their helpful comments. We also thank Bob Schnorbus 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 this paper are those of the authors and do not necessarily reflect the views of members of the Board of Governors or the views of other members of the staff of the Federal Reserve System.

This paper can be downloaded without charge from the Social Science Research Network
electronic library at: <http://ssrn.com/abstract=817366>

Prices, Production and Inventories over the Automotive Model Year*

Adam Copeland
Board of Governors of the Federal Reserve System
adam.copeland@gmail.com

Wendy Dunn
Board of Governors of the Federal Reserve System
wdunn@frb.gov

George J. Hall
Yale University
Department of Economics
george.hall@yale.edu

Abstract

This paper studies the within-model-year pricing and production of new automobiles. Using new monthly data on U.S. transaction prices, we document that for the typical new vehicle, prices typically fall over the model year at a 9.2 percent annual rate. Concurrently, both sales and inventories are hump shaped. To explain these time series, we formulate a market equilibrium model for new automobiles in which inventory and pricing decisions are made simultaneously. On the demand side, we use micro-level data to estimate time-varying aggregate demand curves for each vehicle. On the supply side, we solve a dynamic programming model of an automaker that, while able to produce only one vintage of a product at a time, may accumulate inventories and consequently sell multiple vintages of the same product simultaneously. The profit maximizing pricing and production strategies under a build-to-stock inventory policy imply declining prices and hump-shaped sales and inventories of the magnitudes observed in the data. Further, roughly half of the price decline is driven by inventory control considerations, as opposed to decreasing demand.

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

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

Two common features of durable goods markets are high levels of inventories relative to sales and declining prices over the product cycle. The optimal pricing strategy and inventory management of a good over its product cycle are classic issues in the industrial organization, revenue management, and inventory literatures. However much of the work explicitly linking pricing and inventories policies has been theoretical, while much of the empirical literature has studied either inventories or pricing in isolation. In this paper we jointly consider the optimal pricing, production, and inventory management policies for a durable goods producer and quantify empirically the role inventory control has on the time path of prices within the product cycle.

We accomplish this by analyzing the pricing and production decisions of an automaker over the model year, taking into account both the annual introduction of new vintages and the need to maintain sufficient inventories to facilitate sales. We first document, by vintage, the within-model-year inventory, sales, and pricing behavior in the U.S. market for new automobiles. We then formulate a market equilibrium model for new automobiles. We estimate “typical” within-model-year demand curves for each market segment and vintage and solve a dynamic programming model of the firm with overlapping vintages. The dynamic production and pricing rules implied by the model incorporate not only the sales life cycle of each vintage but also the competition across vintages in the new-vehicle market. We find that the interaction between inventory and pricing policies are key to matching the magnitude of both the observed price declines and the ratio of inventory to sales.

Our data on prices, production, and sales of automobiles over the product cycle are the result of merging two datasets. We match new model-level monthly data on U.S. transaction prices with well-known data on production and sales. These new price data are of unusually high quality for the automobile industry because they not only record the actual transaction price (not the list or invoice price) but also take into account rebates and financing incentives the customer received. Using these data, we document five facts:

1. For the typical new vehicle, the average monthly decline in retail prices (net of rebates and incentives) is 9.2 percent at an annual rate.
2. Although new vintages of a vehicle are introduced annually, the average vintage is sold for 16.7 months. Thus, for nearly half of each calendar year, two vintages of each model are sold simultaneously.
3. When two model years of the same make are selling simultaneously, the old vintage sells for a price 8.8 percent less, on average, than the price of the new vintage.

4. Both sales and inventories are hump shaped over the model year, where the mean ratio of inventories to sales is 76 days.
5. All other things being constant, higher inventories are associated with lower retail prices.

Facts 1 through 3 describe the price path of a typical model within and across model years. 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, microprocessors, and consumer electronics. The sale of several vintages of the same good is a potentially important dimension to the firm's problem, because the within new-market competition is in addition to, and potentially more important than, competition from the used market. New goods of an older vintage are usually quite similar to those of the newer vintage, and they do not suffer from the asymmetric information problems inherent in the used good market. Most existing theories that seek to explain the price paths described by facts 1 through 3 focus on intertemporal price discrimination (e.g. Stokey, 1979) or fashion (e.g. Lazear, 1986; Pashigian, 1988; and Pesendorfer, 1995). These theories, however, place little emphasis on the firm's production decisions. Perhaps not surprisingly, we find that with regard to the market for new automobiles, these theories alone cannot explain the contemporaneous comovements in prices, sales, and inventories over the product cycle.

The last two stylized facts regarding sales and inventories are well known within the automobile industry. Automakers carry an extraordinary amount of inventory relative to sales, as lots at dealerships are kept fully stocked with vehicles. In combination with declining prices over the product cycle, these facts suggest that a model with either stable demand or stable supply will be unable to replicate the first three facts. During the first six months a vehicle is sold, prices are high but declining, quantities sold are low but rising, and inventories accumulate. This pattern suggests that rightward shifts in the supply curve dominate changes in the demand curve early in the model year. During the last twelve months the vehicle is sold, however, both prices and sales fall, an indication that leftward shifts in the demand curve now play the major role.

To refine this intuition, we formulate a market equilibrium model for new cars that links the operations research literature on optimal inventory and revenue management with the economics literature on discrete-choice models of product differentiation. In particular, we solve a dynamic structural model with overlapping vintages in which an automaker can adjust both the price and the quantity produced within the model year. The automaker sells a vehicle that is slightly modified, or changes vintage, every year. While

the automaker produces only the current vintage, the use of inventories allows the firm to sell more than one vintage of the product simultaneously. Each week, the firm must decide the number of units of the current vintage to produce and the optimal prices for the vintages in stock.

Each week the firm faces a different downward sloping demand curve for each vintage. We derive these demand curves from estimates of consumer 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 this standard approach in three main ways. First, we have a better measure of prices, as we use transaction prices instead of the usual list prices. Second, we estimate our demand-side model at a quarterly, rather than an annual, frequency; thus, we estimate how the demand curve shifts throughout the model year. Third, we allow consumers to choose among multiple vintages within and across models. Using our estimates of consumers' preferences, we compute average demand curves for each automobile market segment (such as compact cars) and vintage over the automobile product cycle. An advantage of our approach is that we allow households to differ across quarters in their distaste for price, hence our model accommodates, at least in part, the possibility of the firm engaging in intertemporal price discrimination.

A main result from the demand side analysis is that demand curves for most vehicles shift significantly leftward over the second half of the model year, and the slopes of the curves undergo small changes. Further, we estimate that cross-price elasticities between models of different vintages are quite small. Hence, despite the similarities between vintages of the same model, consumers view vehicles from one model year as poor substitutes for vehicles in another model year. These estimates play a central role in the firm's problem, as they directly affect the firm's revenue flows from selling multiple vintages of the same model at the same time. Lastly, we find that households are slightly more price sensitive during periods of the year when manufacturers typically offer two vintages of a model for sale, relative to those quarters where only one model year is usually available.

Taking the demand curves as given, the firm solves a dynamic inventory problem to maximize profits. 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 to have a cost structure of producing vehicles with several nontrivial nonconvexities.

¹Federgruen and Heching (1999) and Elmaghraby and Keskinocak (2003) provide a nice overview of the more recent revenue management literature within operations research. Also see Chan, Hall, and Rust (2005) for an analysis of a similar pricing/procurement decision.

A significant aspect of the automotive market that we reflect in our model is the distribution of dealerships across the geographic market. Showrooms are instrumental in allowing consumers to learn about manufacturers' products and to gauge products' characteristics. In this industry, for example, consumers value the ability to observe the vehicle they are considering purchasing and to take possession of the vehicle without delay. Consequently, part of the automaker's problem is ensuring that there are a sufficient level of inventories on dealer lots across the national market. The automotive trade press often mentions the necessity of a showroom presence when discussing manufacturers' inventories. In a recent issue of *Ward's Automotive Reports* (August 2, 2004), a Cadillac executive stated, "We have 1,000 dealers who sell less than 50 vehicles a year. They're holding 300 to 400 days' supply [that is, inventories over sales] because they want to display all the models." We incorporate this need to build-to-stock (that is, inventories are a prerequisite for sales) into the model by assuming that the firm faces a "revenue tax" that is a function of sales over inventory. We assume that increases in the sales-to-inventory ratio make it harder for the firm to consummate sales by raising the tax the firm pays per transaction.

After calibrating the supply-side parameters of the model, we are able to replicate the decline in prices over the model year along with hump-shaped sales and inventories that we observe in the data. Early in the model year, the automaker sets the vehicle price high to dampen sales and thus accumulate a large stock of inventories. Building up inventories, or following a build-to-stock inventory management strategy, is optimal because it reduces the cost of carrying out a transaction (that is, it lowers the revenue tax). Over the remainder of the model year, our estimate of leftward-shifting demand lowers the shadow value of inventories, resulting in a 8.1 percent decline in the retail price of a vehicle over the entire product cycle and an average vintage premium of 7.9 percent.

We then solve the model with no revenue tax, thus allowing the firm to manage inventories on a build-to-order basis. In this case, retail prices fall over the product cycle but by only one-half of the magnitude we see in the data. Hence, unlike previous work that attributes falling prices to fashion or price discrimination, our theory implies that build-to-stock inventory management is as important in driving price decline as these other forces. Further, under a build-to-order policy the paths of sales and inventories do not feature the prominent hump-shaped patterns seen in the data. These results demonstrate the significance of the firm's inventory strategy on the optimal pricing path. Falling demand alone explains about half of the price decline in automobiles and misses significant comovements among prices, sales and inventories.

1 Data Sources and Empirical Observations

In this section, we outline the sources of the data used in our analysis and document several stylized facts.

1.1 Data Sources

To construct a dataset with information on 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 provides information on prices, by model and model year, and on the distribution of sales, also by model and model year. The second data source contains information on 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), who obtained the data from J.D. Power and Associates (JDPA). JDPA collects daily transaction-level information from dealerships across the U.S., which it aggregates to a monthly frequency. Then, along the product space dimension, JDPA adds up the data to a model and model-year level. The sample of transactions we use represents 70 percent of the geographical markets in the U.S. and roughly 15 to 20 percent of national retail transactions. It contains monthly observations for almost all unique make, model, and model-year light motor vehicles (for example, 2000 Ford Escort) sold in the U.S. and covers the period from January 1999 to January 2004. Among other variables, the dataset includes information on the number of transactions recorded, the average transaction price, the average cash rebate, and details about the average financial package customers received. JDPA attempts to precisely measure the transaction price of a vehicle. This measure includes the price of accessories (such as roof racks) and transportation costs but excludes aftermarket options, taxes, title fees, and other documentary 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. 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.

In the data, we observe the amount financed, interest rate, and loan term that the average customer received. The financial data are captured for loans that customers obtained from any financial institution, as long as the financing was 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 approxima-

tion 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.² The value of the manufacturer's financial incentive is then defined as the difference between the two NPV amounts. Finally, we convert the market price into 2000 dollars by using the BEA's personal consumption deflator.

As stated earlier, the data from JDPa provide the market prices of vehicles and the distribution of sales, by model and model year. Using the total number of transactions across model years in each period, we compute the fraction of a model's sales that is accounted for by each available model year.

We linked the JDPa data to information on the U.S. sales and North American production of General Motors, Ford, and Chrysler, which we obtained from Ward's Communications. We excluded foreign manufacturers, as measuring overseas production is difficult. The sales data for these firms (also known as 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 on the basis of the monthly model-year distributions in the JDPa sample. Using information from Ward's on model changeover dates at North American assembly plants, we decomposed the production data by model into observations by model year. Finally, using the sales and production estimates by make, model, and model year, we constructed estimates of vehicle 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 market 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 stated earlier, by examining these data we can observe several stylized facts that hold across models and model years. To provide illustrative examples, we show plots of the price, sales, production, and inventory data for a midsize car and a pickup truck (figures 1-8). The steady decrease in price over the sales cycle is immediately evident for both vehicles shown in the figures. In the 2000 model year, the average market 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

²The Board of Governors publishes these data in its G.19 Consumer Credit release

Market Segment	Model Year					All	
	1999	2000	2001	2002	2003		
Compact	7.7	5.9	8.1	9.4	17.5	9.5	(2.4)
Midsized	9.1	6.7	6.1	9.0	16.4	9.2	(1.5)
Fullsize	8.9	7.9	6.4	8.5	13.4	8.9	(2.1)
Luxury	11.6	10.3	8.8	13.1	14.9	11.6	(1.2)
Pickup	6.6	10.0	7.1	9.2	15.4	9.9	(2.2)
SUV	7.0	6.7	7.2	5.2	13.6	8.2	(0.9)
Sporty	2.3	6.2	0.4	6.1	10.9	5.1	(2.4)
Vans	5.4	9.0	9.3	8.5	15.7	9.6	(1.4)
Total	7.4	8.0	7.1	8.1	15.1	9.2	(0.6)

Note: Standard errors are in parenthesis

Table 1: The Average Monthly Price Decline (annual rate) by Market Segment and Model Year

\$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 the midsize car is almost 7 percent, while the premium for the pickup averages about 9 percent. Turning to the figures on sales and inventories, we find that the sales and inventories of both the midsize car and the pickup exhibit a hump-shaped profile.

Table 1 provides a summary of the average monthly price decline, the first stylized fact, by market segment and model year and weighted by sales. For the midsize market segment, the mean monthly decline in prices 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. 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. This reflects especially high incentives offered by manufacturers in the latter half of the product cycle. Overall, the monthly decline in price averages 9.2 percent at an annual rate.

To observe the within-year price declines more generally, we illustrate the aggregate matched-model price indexes for successive model years as constructed by Corrado, Dunn, and Otoo (2004) (figure 9). This price index was constructed from the entire JDPA dataset and so includes price data on vehicles produced by European and Asian automakers. As can be seen, transaction prices for a given model year are at their highest levels when each model is introduced, and they trend downward in a consistent pattern over the course of the sales cycle. The overlap of the various model-year price indexes highlights the

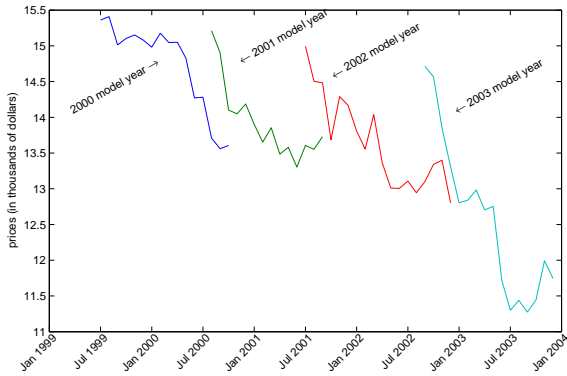


Figure 1: Average Transaction Prices.




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.

second stylized fact—that multiple vintages of vehicles are simultaneously sold for a significant portion of the model year. In our database of transactions, the mean length of time a vehicle is on the market is 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 the average new vintage premium by market segment and model year, weighted by sales (table 2). The table illustrates that the new vintage premium of 2000 model-year midsize cars over their 1999 model-year counterparts is 10.0 percent. In the sample, the new vintage premium is 8.5 percent, on average, for midsize cars, and the standard error is 0.4.

Although the new vintage premium varies quite a bit across market segments and time, overall it amounts to 8.8 percent, on average, in our sample. This premium is highest for luxury cars and pickup trucks and lowest for compact cars, sport utility vehicles, and sporty cars, where the difference in premiums between luxury and compacts is 4.5 percent. Across model years, the average new vintage premium is typically between 7 and 9 percent, though the premium during the 2003 to 2004 changeover is 13.1 percent. This large premium is related to the steep decline in within model-year prices for 2003 model-year vehicles, shown in table 1. Both numbers reflect the unusually high incentives the Big Three placed on 2003 model-year vehicles over the second half of the product cycle.

One might argue that the new vintage premium simply reflects improvements in quality or additional

Market Segment	Model Year					All	
	2000	2001	2002	2003	2004		
Compact	5.9	6.9	6.7	6.9	11.0	7.1	(0.5)
Midsized	10.0	5.8	6.1	7.5	11.9	8.5	(0.4)
Fullsize	9.9	6.2	7.8	8.2	9.2	8.3	(0.6)
Luxury	11.0	11.6	9.8	14.1	11.1	11.6	(0.4)
Pickup	10.7	9.8	6.3	8.6	20.0	10.6	(0.7)
SUV	5.4	0.4	10.1	8.8	10.9	7.2	(0.4)
Sporty	2.6	7.5	3.3	28.9	-7.8	7.2	(0.8)
Van	7.7	11.8	3.6	9.1	12.4	8.6	(0.4)
All	8.6	7.2	7.1	9.1	13.1	8.8	(0.2)

Note: Standard errors are in parenthesis

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

features. For example, the 20 percent new vintage premium recorded for 2004 model-year pickup trucks reflects, in part, a quality improvement made to Ford’s F-series pickup truck.³ For many 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 this fact, we looked at the new vintage premium for a subsample of vehicles that had not undergone a major redesign. We determined when a model received a major redesign by using data from Ward’s Communications on the vehicle’s platform. Given that the platform choice designates the basic structure of the vehicle, we take a manufacturer’s decision to change a vehicle’s platform as a sufficient, though not necessary, condition that the vehicle has undergone a major redesign. We recomputed the new vintage premium for these vehicles and found that the average premium differs little from the figures reported in table 2.

The fourth stylized fact states that both sales and inventories for a particular model and model year exhibit a humped-shaped pattern. When we plot the Big Three’s aggregate sales by model year, we see the distinctive hump shape that sales follow over the product cycle (figure 10). The contour of aggregate sales, however, confounds the evolution of sales over the product cycle with calendar effects, because vehicles of a given model year are not all introduced in the same month. To separate out these two effects, we define the dummy variables 1_t for $t = 1, 2, \dots, 14$ as indicators of how many months the model year has been sold. If $1_{t=1}$ is equal to 1, then at this date the associated vehicle is in its first month of sales. For model

³Alternatively, if a cheaper base model is introduced, the vehicle premium may be biased downwards (see, for example, the -7.8 premium for 2004 sporty cars).




Figure 10: Big Three Sales by Model Year

years that are sold for more than fourteen months, we define a dummy variable 1_{15+} , which is equal to 1 for sales that take place more than fourteen months after the introduction of the vehicle to the marketplace. We then run a regression of log sales on these dummy variables, including fixed effects for each model and controlling for the calendar month. Examining the coefficients on these dummy variables then provides an estimate of the shape of the typical sales path for a vehicle. Using a similar approach, but substituting inventories for sales, we can also study the shape of the typical inventory path for a vehicle. We report the estimated coefficients from running these two regressions (table 3).

Regarding the sales regression, the estimated coefficients imply that sales do indeed follow a hump-shaped path. After rapid sales growth in the first three selling months, the level of sales then slowly climbs to its peak six to eight months after the vehicle's debut before declining. This same pattern can be seen in the estimated coefficients from the regression using inventories. Inventories peak slightly later, roughly nine months after a vehicle's debut, before rapidly falling.

To better analyze the relationship between sales and inventories, we consider the ratio of inventories to sales, also known as days' supply. We use the stock of inventories at the beginning of the period and sales for the current month. Hence, 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, assuming future sales are equal to the current month's flow. A major focus of the days' supply calculation is the large stock of

Parameter	Sales		Inventories	
	Coefficient	Standard Error	Coefficient	Standard Error
$1_{t=2}$	0.76	0.038	0.32	0.058
$1_{t=3}$	1.09	0.039	0.50	0.060
$1_{t=4}$	1.24	0.039	0.62	0.061
$1_{t=5}$	1.24	0.039	0.79	0.061
$1_{t=6}$	1.25	0.040	0.94	0.062
$1_{t=7}$	1.26	0.041	1.03	0.066
$1_{t=8}$	1.25	0.041	1.05	0.067
$1_{t=9}$	1.18	0.041	1.11	0.068
$1_{t=10}$	0.98	0.042	1.09	0.067
$1_{t=11}$	0.93	0.042	0.84	0.066
$1_{t=12}$	0.95	0.042	0.38	0.064
$1_{t=13}$	0.82	0.041	-0.07	0.063
$1_{t=14}$	0.60	0.043	-0.45	0.066
$1_{t=15+}$	-0.11	0.034	-0.80	0.047

Table 3: The Shape of Sales and Inventories by Model and Model Year Over the Product Cycle

Market Segment	Compact	Midsize	Fullsize	Luxury	Pickup	SUV	Sporty	Vans	All
Days Supply	73	60	75	80	84	75	83	85	76 (2.1)

Note: Standard errors are in parenthesis

Table 4: The Average Days Supply by Market Segment

inventories that automakers carry relative to sales. On average, automakers carry 76 days' supply in our data—implying automakers typically carry enough inventories to sell vehicles for over two 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.

Turning to the last stylized fact, we look at the correlation between inventories and prices, by model and model year. To analyze the relationship between prices and moments when inventories are above or below trend, we first need to accurately measure when inventories are ample or lean. The residuals from the inventory regression described earlier provide a measure of the deviations from the usual contour of inventories over the product cycle. We measure whether prices are correlated with these inventory fluctuations by regressing the log of price on a lag of these inventory residuals. We find the expected significant negative relationship between lagged inventory residuals and price: The estimated coefficient on the lagged residuals is -0.02, and the associated standard error is 0.003.

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 the new vintage premium reflects optimal pricing behavior in an environment in which demand is driven by fashion. They use monthly data from the consumer price index to show that at the aggregate level prices for new cars decline between December and September of the model year. Although the magnitude of the within-model-year price declines have fallen between 1954 and 1989, Pashigian, Bowen, and Gould find these price declines are larger for luxury and speciality cars than for compacts and subcompacts. They argue that the larger the changes in styling and quality improvements between model years, the larger the within-model-year price declines. Hence, the new-car market behaves like the market for fashionable apparel. In contrast, we find only marginal evidence of a particular pattern in the new vintage premium across market segments. As shown in table 2, although luxury cars command the highest vintage premium, other fashion-oriented vehicle types such as SUVs and sporty automobiles command premiums nearly equal to, or markedly below, those of the more plain compact and midsize automobiles.

Given the evidence, we posit that within-model-year price declines are driven more by the used-vehicle market than by fashion. Consider the case of a 2000 model-year vehicle produced at the end of the product cycle and a 2001 model-year vehicle produced at the beginning of the product cycle. Although these two vehicles may have been produced just a few weeks apart, we expect that in the used-vehicle market, the 2000 vintage will be perceived to have been on the road twelve months longer than the 2001 vintage.

To provide evidence in support of this hypothesis, we estimated 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 in the regression 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.⁴ We show the resulting coefficients on age and odometer reading, both of which are statistically significant with greater than 99 percent confidence (table 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 on a given vehicle. Notably, the coefficient on age implies that, even after controlling for the odometer reading and other vehicle characteristics, a higher model age (that is, an older model year) implies a lower price in the used vehicle market. All other things constant, increasing the age (as defined by the model year) by one

⁴Because 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.

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

Table 5: Coefficients on Age and Odometer from the Used-vehicle Price Regression

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.

2 A Market Equilibrium Model with Overlapping Model Years

In this section, we present a market equilibrium model designed to capture the empirical regularities documented earlier. 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 the automaker is able to sell two vehicles, the current year's vintage and the previous year's vintage. The firm's model parameters are calibrated to match the key features of the firm's cost structure and the means of prices, output and inventories. We derive decision rules that govern the production and pricing of vehicles over the model year. Through numerical simulations, we demonstrate that the empirical regularities documented earlier are consistent with our derived decision rules under a build-to-stock inventory policy.

An essential feature of the firm's problem is the market demand curve. We posit a semi-log demand curve whose parameters are price semi-elasticities. We then draw upon the existing demand-choice literature to estimate these semi-elasticities and their change over the product cycle.

2.1 The Automaker's Problem

In the interest of tractability, we make several strong simplifying assumptions on the supply side. First, we assume that 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 issues of bargaining and price discrimination 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 actually decentralize their operations both across products

and between the production and marketing sides of the business. But because these issues are not central to understanding the facts presented earlier, we defer further consideration to other papers.⁵

The automaker sells two products: this year's vintage and last year's vintage. The decision period is a week. There are T 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 T -week model-year problem. Successive model years are linked because this year's vintage becomes last year's vintage at the end of the T^{th} 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^{this} ; 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 and the price of the other vintage:

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

where μ_t^j is a constant term, η_t^j is the own-price semi-elasticity, and ϕ_t^{ji} is the cross-price semi-elasticity. The demand parameters may vary across the 52 weeks of the year.

Unsold vehicles can be inventoried without depreciation. Let I_{t+1}^j be the stock of vintage j vehicles that are inventoried at the end of period t and carried over into period $t + 1$. Current production is not available for immediate sale, so sales can be made only from the beginning-of-period inventories:

$$s_t^j \leq I_t^j. \quad (2)$$

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

$$I_{t+1}^{this} = I_t^{this} + q_t - s_t^{this}. \quad (3)$$

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

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

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:

$$I_1^{last} = I_T^{this} + q_T - s_T^{this}. \quad (5)$$

⁵For example, Bresnahan and Reiss (1985) model and estimate the division of markups between automobile manufacturers and dealers. For discussions of bargaining and price discrimination in the retail auto market see Ayres and Siegelman (1995), Goldberg (1996), and Zettelmeyer, Scott Morton, and Silva-Risso (2001).

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 linear in hours:

$$q_t = D_t \times S_t \times h_t \times LS. \quad (6)$$

Although the production function is linear, the firm faces several important non-convexities because of its labor contract.⁶ The average straight-time, day-shift wage at these plants is about \$27 an hour plus benefits. Workers on the second (evening) shift receive a 5 percent premium. 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 a worker receives for the 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

$$\begin{aligned} c(D_t, S_t, h_t | q_t) = & \gamma q_t + (w_1 + I(S_t = 2)w_2) \times (D_t h_t n + \max[0, 0.85(40 - D_t h_t)n]) \\ & + \max[0, 0.5D_t(h_t - 8)n] + \max[0, 0.5(D_t - 5)8n] + 0.65w_1 40(2 - S_t)n, \end{aligned} \quad (7)$$

where γ is the per vehicle material cost, 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. The first term is the per-vehicle cost; it 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. The subsequent terms within the brackets capture the 85 percent rule for short weeks and the required overtime premium. The last term is the unemployment compensation

⁶For further discussion of the institutional details of labor contracts in automobile manufacturing, see Bresnahan and Ramey (1994), Hall (2000), or Ramey and Vine (2004).

bill charged to the firm. Let $D_t = 0$ if and only if $S_t = 0$. Thus, the 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 will minimize average costs by operating the plant with either one shift or two shifts for 40 hours per 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

$$\sum_{t=1}^T \left(\frac{1}{1+r} \right)^{t-1} \left\{ p_t^{last} s_t^{last} (1 - \tau(s_t^{last}/i_t^{last})^\Psi) + p_t^{this} s_t^{this} (1 - \tau(s_t^{this}/i_t^{this})^\Psi) - c(D_t, S_t, h_t | q_t) \right\} + \left(\frac{1}{1+r} \right)^T V(I_{T+1}^{last}, 0, 1) \quad (8)$$

subject to (1)-(6) and where $c(D, S, h | q)$ is given by (7). The terms $\tau(s_t^{last}/i_t^{last})^\Psi$ and $\tau(s_t^{this}/i_t^{this})^\Psi$ are revenue taxes the automaker must pay if the sales-to-inventory ratio is large. This term captures the distributional costs the automaker faces, as described previously in the introduction. When inventories are low, it is harder for potential customers to observe and gauge a vehicle (that is, to test-drive it and view the choice set), and thus it is more costly to consummate a sale. The tax effectively disappears when the sales-to-inventory ratio is small.⁷ The term $V(I_{T+1}^{last}, 0, 1)$ 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 can be written:

$$V(I^{last}, I^{this}, t) = \max_{p^{this}, p^{last}, q} \left\{ p^{last} s^{last} (1 - \tau(s^{last}/i^{last})^\Psi) + p^{this} s^{this} (1 - \tau(s^{this}/i^{this})^\Psi) - \min_{D, S, h} c(D, S, h | q) + \frac{1}{1+r} V(I^{last} - s^{last}, I^{this} + q - s^{this}, t+1) \right\} \quad \text{for } t = 1, \dots, T-1 \quad (9)$$

subject to (1), (2), and (6) and where $c(D, S, h | q)$ is given by (7).

At week T , this year's vintage becomes last year's vintage, and so the value function is

$$V(I^{last}, I^{this}, T) = \max_{p^{this}, p^{last}, q} \left\{ p^{last} s^{last} (1 - \tau(s^{last}/i^{last})^\Psi) + p^{this} s^{this} (1 - \tau(s^{this}/i^{this})^\Psi) - \min_{D, S, h} c(D, S, h | q) + \frac{1}{1+r} V(I^{this} + q - s^{this}, 0, 1) \right\}. \quad (10)$$

⁷We tried specifying demand as an increasing function of the level of inventories; that is, customers are willing to pay a higher price if they are more likely to be matched to their ideal vehicle. However, in this case, the model has the counterfactual implication that, all other things being equal, higher inventories are associated with higher prices. Fact 5 suggests that, even if demand is increasing in inventories, the supply-side inventory effect dominates the demand-side effect. We also tried replacing the revenue tax with an explicit inventory-to-sales target; this yields similar results, but the tax is computationally more robust.

2.2 Parameterizing the Model

There a number of parameters in this model. Our approach is to choose the supply-side parameters in the firm's problem based on the first moments in the data and from published information on assembly plants. For demand parameters, we employ a discrete-choice methodology to estimate consumer's preferences over automobiles. We then use these estimates to compute the intercepts, own-price semi-elasticities and cross-price semi-elasticities that are parameters in the market demand function, equation (1).

2.2.1 Demand-side parameters

Overview: The demand for automobiles is modelled within a discrete-choice framework. 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 horsepower and miles per gallon), which are denoted by X ; and those that are unobserved by the econometrician (such as styling or prestige), which are denoted by ξ . Drawing from the nested logit literature, we also incorporate a correlation in the consumer's tastes for vehicles of the same model year. We divide vehicles into $G + 1$ mutually exclusive groups (that is, model years)— $g = 0, 1, 2, \dots, G$ —where the outside good is the sole member of group 0. We also allow households' distaste for price, denoted by α , to vary from quarter to quarter. This captures 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 , dropping the time subscript, as

$$u_{ijq} = X_j\beta + \xi_j - \alpha_{iq}p_j + \sum_k \sigma_k v_{ik} x_{jk} + \zeta_{ig} + (1 - \rho)\varepsilon_{ij}, \quad (11)$$

where p_j denotes the price of product j and $x_{jk} \in X_j$ is the k th observable characteristic of product j . The term $X_j\beta + \xi_j$, 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. 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 increases in quarter $q = \{1, 2, 3, 4\}$. 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. We assume the distribution of household income is lognormal, and, for each year in our sample, we estimate its mean and variance from the Current Population Survey (CPS). The second-to-last term in equation (11) captures correlations in a consumer's tastes for products within the same group. For consumer i , the variable ζ_{ig} is common to all products in group g and has a distribution that depends upon ρ . Finally, ε_{ij} is an i.i.d. extreme value.

Consumers choose among the $j = 1, 2, \dots, J$ automobiles in our sample and the outside good (denoted $j = 0$), which represents the choice not to buy a new automobile from the Big Three. Consumers choose the product j that maximizes utility, and market shares are obtained by aggregating over consumers.

Implementation: As described in section 1, our sample includes data for 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. We aggregate sales and prices to the quarterly frequency because there is a significant amount of volatility in monthly sales due, in part, to intertemporal substitution. Moreover, BLP's static utility maximization approach is better suited to analyzing quarterly data. We do not estimate the model at an annual frequency because the variation in price and in the consumer's choice set from quarter-to-quarter is a significant source of identification in the BLP framework.

We use vehicle characteristics that include a set of model dummies, a measure of acceleration, vehicle dimensions, a measure of safety, and fuel efficiency.⁸ As was done in previous research, we link quantity sold and transaction price to the characteristics of the base model to produce a vehicle-quarter observation. To this standard set of characteristics, we also add a measure of how long a product has been sold—the number of quarters since the vehicle was first introduced. The model dummies, the number of quarters since introduction, and a quadratic time trend make up the vector of observable characteristic used to compute the mean utility of a product, $X_j\beta$. The measures of acceleration, dimension, safety, and fuel efficiency, along with the number of quarters since introduction, are included in the vector of observable characteristics used to measure heterogeneity in households' preferences, $\sum_k \sigma_k v_{ik} X_{jk}$. In essence, the model dummies help explain the mean utility level of a product, while the measures of acceleration, dimension, safety, fuel efficiency, and number of quarters since introduction drive the substitution patterns among vehicles.

Unlike the approach taken in BLP, we also incorporate (through ζ) correlation in consumers' tastes across a discrete characteristic, model years. This nested logit approach is folded into the BLP algorithm in the natural way. Following BLP, we use the number of households in the U.S. as reported in the CPS as

⁸Information on vehicle characteristics were taken from Automotive News's *Market Data Book* (various years).

a measure of market size for the year. We assume that one-fourth of all households in a given year show up each quarter.

Our estimation strategy follows the generalized method of moments approach taken by BLP. Given the vector of parameters θ , we solve for the unique vector of mean utilities such that the model's predicted market shares equal actual market shares.⁹ We then match the moments related to the market-level disturbance, ξ_j , using the assumption that ξ is uncorrelated with the vehicle characteristics, X , or

$$E[\xi(\theta)|X] = 0. \quad (12)$$

As in the typical nested logit exercise, the parameter ρ is estimated through the decomposition of the mean utility of a product,

$$\delta_j = X_j\beta - \xi_j + \rho \ln(s_{j/g}), \quad (13)$$

where δ_j is equal to the difference between the market share of product j and the market share of the outside good, and $s_{j/g}$ is the market share of product j relative to the market share of the group to which product j belongs.

As ξ is correlated with both price and $s_{j/g}$, an endogeneity problem arises. Berry (1994) provides a methodology that allows us to use instrumental variables. We follow BLP's approximation of the optimal instruments, though in our setting they have a diminished effect. These instruments are based on competing products' characteristics, which change at the model-year frequency, though our price and quantity data vary at the quarterly level. Accordingly, we augment the set of instruments to include indicator variables for whether multiple vintages of a model are being sold simultaneously. Whether or not another vintage is being sold at the same time has an effect on price, and this effect can vary at a quarterly frequency. Further, there is little reason to suspect that the sale of multiple vintages of a model is related to the unobserved characteristic of a vehicle. New vintages are most often introduced at an annual frequency, and automakers face large costs to altering the scheduled introduction of a new vintage.

Results: We present a subset of the parameter estimates in table 6. Given their large number, we do not report our fixed-effects estimates. Instead, we show those estimates that measure the heterogeneity in consumers' tastes (σ) along with estimates of the substitutability of vehicles across model years (ρ) and estimates of a consumer's distaste for price (α). The coefficients on miles per dollar (measure of

⁹We 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, which is used to integrate over v to obtain a prediction of market share.

Parameters		Coefficient	Standard Error
Taste for Variety	σ		
	Constant	0.61	0.968
	Horsepower/Weight	3.78	0.905
	Height	3.91	1.086
	Size	2.07	0.752
	Miles/dollar	0.50	2.806
	Air bag	0.93	0.631
	Intro	-0.04	0.154
Model-Year Nest	ρ	0.58	0.014
Distaste for Price (Q1)	α_1	42.31	7.684
Distaste for Price (Q2)	α_2	36.64	7.998
Distaste for Price (Q3)	α_3	34.99	6.823
Distaste for Price (Q4)	α_4	32.95	4.707

Table 6: Parameter Estimates

fuel efficiency), air bag (measure of safety), and intro (number of quarters since introduction) are not statistically significant. However, we estimate that consumers are quite heterogeneous in their tastes for acceleration (horsepower over weight), height, and size (length times width of the vehicle).

The model-year nest and the price coefficients are precisely estimated. The estimate of ρ reflects the substitutability of models from different model years, where $\rho = 0$ indicates no substitutability and $\rho = 1$ implies perfect substitutability. Our estimate of ρ is 0.58, which suggests that consumers do not find vehicles from different model years to be close substitutes. The estimated value of household's distaste for price falls from quarter 1 to quarter 4. 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. Although the differences among the α 's are not significant, they imply that households are more sensitive to price in quarters when automakers typically offer multiple vintages of vehicles, relative to quarters when only one vintage is available.

The estimate of ρ and the magnitude of the price coefficients are more easily interpreted by examining the implied own-price and cross-price elasticities. These elasticities provide the clearest picture of the values of the semi-elasticities that we use in our specified demand function. We report the own-price elasticities of individual vehicles averaged across market segments, quarters, and vintages, where the vintage label signifies whether the vehicle is the newest model year available or not (table 7).

Vintage	Market Segment	1st Quarter	2nd Quarter	3rd Quarter	4th Quarter
New	Compact	8.2	9.2	7.9	8.6
	Full	10.3	11.3	9.1	9.6
	Luxury	9.8	11.4	8.4	9.0
	Midsize	9.4	10.5	8.9	9.7
	Pickup	9.7	10.8	9.1	9.0
	SUV	9.7	10.7	8.6	8.6
	Sporty	10.8	10.9	9.1	10.0
	Van	10.2	11.6	9.7	10.1
	All	9.8	10.8	8.8	9.3
Old	Compact	7.8	8.3	8.8	8.8
	Full	10.0	11.4	10.5	10.6
	Luxury	9.9	11.0	9.2	9.2
	Midsize	9.4	10.0	8.9	8.1
	Pickup	9.8	10.8	8.9	11.6
	SUV	9.9	10.8	9.6	8.6
	Sporty	9.8	11.7	9.3	8.6
	Van	10.2	10.9	9.6	5.8
	All	9.6	10.6	9.3	8.9

Table 7: The Absolute Value of Own-Price Elasticities by Market Segment, Quarter, and Vintage

The own-price elasticities generated by our parameter estimates range between 6 and 12, an indication that manufacturers face quite 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 an 8.2 percent fall in sales, holding everything else equal. The average own-price elasticity across all vehicles is reported in the “All” row, and illustrates that elasticities in the first and second quarter are typically higher than those in the third and fourth quarter for both new and old vintages. This is likely driven by the higher α 's and larger choice sets in these quarters. In general, our estimated elasticities are higher than those found in the previous literature; BLP, for example, report a range of elasticities between 3 and 6. It is not surprising, however, that our elasticity estimates are higher than some observed elsewhere because previous research estimated own-price elasticities among models—that is, at a level of aggregation higher than that of our data. Indeed, when we re-estimate the parameters from data aggregated to the model level, the implied own-price elasticities fall within the range of those reported in BLP.

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 to the firm. We

Vintage	Market Segment	1st Quarter	2nd Quarter	3rd Quarter	4th Quarter
New to Old	Compact	0.02	0.01	0.01	0.04
	Full	0.02	0.00	0.00	0.03
	Luxury	0.02	0.01	0.00	0.01
	Midsized	0.02	0.01	0.00	0.01
	Pickup	0.10	0.02	0.03	0.12
	SUV	0.04	0.01	0.02	0.04
	Sporty	0.02	0.01	0.01	0.01
	Van	0.02	0.01	0.00	0.00
Old to New	Compact	0.01	0.03	0.03	0.01
	Full	0.01	0.02	0.00	0.02
	Luxury	0.01	0.03	0.01	0.00
	Midsized	0.01	0.03	0.01	0.00
	Pickup	0.06	0.18	0.02	0.05
	SUV	0.02	0.05	0.02	0.05
	Sporty	0.01	0.03	0.03	0.00
	Van	0.01	0.04	0.05	0.01

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

report our estimates of the cross-price elasticities between two vintages of the same model, averaged over quarters and across market segments (table 8).

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). However, a fair number of vehicles are introduced at other times in the year, and so we can compute cross-price elasticities throughout the year. The upper portion of table 8 displays 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. The bottom portion of the table shows the opposite relationship—the percentage change in market share for the old vintage given a percentage change in the price of the newer vintage. Generally, the estimated cross-price elasticities are quite small relative to the own-price elasticities. This implies that various vintages of the same model are typically quite imperfect substitutes.¹⁰

To check the robustness of these results, we estimated several different specifications of the consumer’s

¹⁰Ana 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.

utility function. We rejected a specification of utility for which groups are defined at the model level (that is, the correlation in consumer tastes is across all available vintages of a particular model) because ρ was estimated to be greater than 1.¹¹ In addition, both the high own-price and the small cross-price estimated elasticities are consistent across several different utility specifications.¹²

2.2.2 Supply-side parameters

For the parameters in the firm's problem, we set T , the number of weeks in a model year, to 52 and the time-invariant interest rate such that $(1+r)^{-52} = 0.95$. The interest rate is the only cost of holding inventories in the model. To parameterize the cost function, we set the line speed, workers per shift, and wage rates to values typically observed at assembly plants. We set the remaining three parameters, γ , τ and ψ to match the average retail price, average rate of production, and average days-supply of inventories observed in the data. Although we would have preferred to estimate these parameters econometrically, computational issues made such estimation infeasible. The line speed at most North American assembly plants is set between 40 and 60 cars per hour; thus, we fix the line speed to 50 cars per hour. Using the employment data from Hall (2000), we set n to 1300 workers per shift, so the firm employs 2600 workers. We read the wages off the union contract: $w_1 = \$27.00$ per hour, and $w_2 = \$28.35$ per hour. We set γ , the per vehicle cost, to the average retail price observed minus \$1500; γ effectively scales the cost function linearly, and thus we allow γ to differ across market segments. We set the revenue tax parameters τ to 1 and ψ to 1.75 such that we match the average days-supply of inventories observed in the data.

2.3 Model Results

Using these parameter values, we solve the dynamic program given by (9) and (10) via an algorithm suggested by John Rust. Specifically, we merge the T value functions into a single time-invariant Bellman equation:

$$V(I^{last}, 0, 1) = \max_{\{p_t^{this}, p_t^{last}, q_t, D_t, S_t, h_t\}} \left\{ \sum_{t=1}^T \left(\frac{1}{1+r} \right)^{t-1} \left(p_t^{last} s_t^{last} (1 - \tau(s^{last}/i^{last})^\psi) \right. \right. \\ \left. \left. + p_t^{this} s_t^{this} (1 - \tau(s^{this}/i^{this})^\psi) - c(D_t, S_t, h_t | q_t) \right) + \left(\frac{1}{1+r} \right)^T V(I_T^{this} + q_T - s_T^{this}, 0, 1) \right\}.$$

¹¹To be consistent with the hypothesis of random utility maximization, it must be that $\rho \in [0, 1]$.

¹²In general, we were unable to estimate alternate utility specifications that had more parameters than did the specification described in the paper (such as random coefficients on a set of model dummies). These alternative specifications demanded too much of the data.

Market Segment	Data			$\tau = 1$			$\tau = 0$		
	Product	Price	D-S	Product	Price	D-S	Product	Price	D-S
Compact	14,524	\$13,622	73	16,208	\$13,594	70	16,000	\$13,642	12
Midsize	10,886	19,193	60	12,000	19,273	75	16,333	19,075	11
Fullsize	7,184	23,772	75	6,167	24,107	77	9,833	23,708	13
Luxury	3,106	36,032	80	2,167	36,430	80	2,479	36,765	9
Pickup	35,114	23,662	84	17,812	24,121	67	22,900	24,224	9
SUV	11,532	28,660	75	9,117	29,140	74	14,500	28,741	11
Sporty	5,721	27,227	83	3,167	27,331	82	6,500	27,081	21
Van	8,396	22,716	85	9,667	22,868	80	13,667	22,581	10
Average			76			76			12

Table 9: Average Monthly Production, Retail Prices, and Days-Supply by Market Segment

To solve for the fixed point, we carried out the following steps: (1) Guess an initial value for $V(I^{last}, 0, 1)$; (2) solve the T Bellman equations in (9) and (10) through backward recursions; (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.

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 26 points from 0 to 60,000. 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 676 inventory pairs, we maximize the right hand side of equations (9) and (10) over each sales price and level of output. Points off the two inventory grids are approximated using bi-linear interpolation. The two sales prices, p^{last} and p^{this} , may take on any positive value such that quantity demanded remains positive. Finally, we impose only mild seasonality on production, assuming that the plant closes for two weeks in July (weeks 51 and 52) for a model changeover.

In table 9 we report the average monthly rate of production, retail price, and days-supply of inventories for a typical vehicle in each market segment. We then report the corresponding averages implied by the model for the case with the revenue tax ($\tau = 1$) and the case without the revenue tax ($\tau = 0$). In general, the model replicates the average price and quantities produced for each market segment. This should not

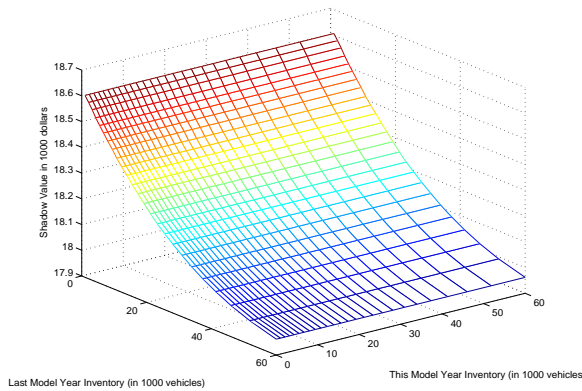


Figure 11: Week 27 Shadow Value of Inventories for This Year's Vintage.




Figure 12: Week 26 Optimal Pricing Rule for This Year's Vintage

be terribly surprising because, given the estimated demand curves, the parameter, γ , was set to match these moments.¹³ With the revenue tax, the model matches the average observed day-supply. Without the revenue tax, however, the average level of inventories are 1/6 the level observed in the data. As is well-understood in the inventory literature, it is difficult to match the high level of inventories observed in many industries without an ad hoc inventory-to-sales target in the firm's objective function.

To illustrate the dynamics of the model, we first consider the firm's pricing decision for a typical midsize car setting $\tau = 1$. Later we report results for which we shut down the revenue tax by setting $\tau = 0$. In figure 11, we plot the partial derivative of the value function for week 27 with respect to inventories of the current model year (other weeks are qualitatively similar). This figure illustrates the shadow value of inventories, or marginal increase in the firm's net worth from an additional unit of inventory, at each point in the state space for week 27. The shadow value of inventories is a decreasing function of the level of inventories. When the level of inventory for this year's model is close to zero, an additional vehicle of this year's model is worth \$18,600 to the firm; however, at the upper bound of the inventory grid, the shadow value of an additional vehicle from this model year is worth only about \$17,960 to the firm—roughly the average cost of producing a vehicle when the plant operates two forty-hour shifts per week. Given the non-convexities in the cost function, computing and reporting the marginal cost of an additional vehicle is bit more involved, but the value for γ (the material cost per vehicle) of \$17,693 (\$19,193-\$1,500) provides a lower bound on the marginal cost.

¹³The model has difficulty matching the quantities of pickup trucks. While we assume each vehicle is made at a single plant, several popular pickup trucks (e.g. Ford F-series, Chevy Silverado, and Dodge Ram) are produced at four or five plants.

Because the automaker faces a downward-sloping demand curve, the profit-maximizing price sets marginal revenue equal to the shadow value of inventories next period. If we set the cross-price semi-elasticities equal to zero and ignore the revenue tax, the optimal price for this year's model is

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

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.

We then plot the pricing rule for this year's vintage for week 26 in figure 12. The pricing rule is almost the shape of the shadow value of inventories. Holding all other things constant, the optimal price is a decreasing function of the level of inventory—our fifth fact. In figure 13, we plot slices of the pricing rules for different weeks in the product cycle, holding the inventory of the competing vintage fixed. Prices are a decreasing function of the level of inventory and the pricing curve shifts down over time. These price rules are consistent with the findings of Zettelmeyer, Scott Morton, and Silva-Risso (2003) that the average retail price at a dealership with ample inventory is about \$230 per car less than that at a dealership with low inventory. These pricing rules do not guarantee that prices fall over the product cycle; after week 52, for example, inventories are monotonically decreasing over time (since there is no replenishment), therefore prices may go up or down depending on the evolution of inventories.

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 designed to be 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.

The implied time series from the model are consistent with the five facts put forth in the introduction. As these figures illustrate, the model generates both downward sloping price paths and hump-shaped inventory and sales. The revenue tax term plays a key role in this. Early on in the model year, inventories are naturally low, so it is expensive to sell a lot of vehicles. In order to reduce this tax in the future, the automaker needs to build up inventories. Hence the automaker sets prices high early on in the year to dampen down sales and allow inventories to accumulate. Once inventories are high (inventories peak in the seven month of the model year), the tax effectively disappears and the firm lowers prices in order to stimulate sales. Further exacerbating the fall in prices, demand for the vehicle starts to decrease as the model year progresses. For forty-one weeks (about three-quarters of the year) the automaker sell both




Figure 13: Optimal Price Rule Over the Product Cycle, Holding the Inventory Level of the Competing Vintage Fixed.

vintages simultaneously and prices fall from \$20,061 to \$17,470: a drop of 12.9 percent over the product cycle or an average monthly price decline of 8.4 percent at an annual rate. This results in a new vintage premium of 8.1 percent, which is within a single standard error of the 8.5 percent new vintage premium for a typical midsize car, as reported in tables 2 and 10.

More generally, as reported in table 10, the model with τ set to 1 is able to match the observed price declines and vintage premia for all market segments. For seven of the eight market segments, the implied price declines are within a single standard error of the average declines seen in the data; and the remaining market segment (luxury cars) is well within the two-standard error band. The model underestimates the average vintage premia by nine-tenths of a percent (8.8 versus 7.9). While this is outside the two-standard error band, the model gets the magnitude right. We believe that relaxing our assumption that new vintages arrive strictly every 52 weeks would enable the model to better match this moment.

For most of the year in the midsize car case, the automaker produces 4000 vehicles per week running two eight-hour shifts for five days per week. For the last two months of the model year, the plant runs just a single shift before shutting down and changing vintages. Comparing figures 3 and 7 to figure 16, it is clear the model predicts production to be quite smooth relative to what is observed in the data. At most




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: $\tau=1$.

Market Segment	Data				$\tau = 1$		$\tau = 0$	
	Price Decline		Vintage Premium		Price Decline	Vintage Premium	Price Decline	Vintage Premium
Compact	9.5	(2.4)	7.1	(0.5)	10.0	9.3	4.6	5.0
Midsize	9.2	(1.5)	8.5	(0.4)	8.4	8.1	3.1	4.3
Fullsize	8.9	(2.1)	8.3	(0.6)	7.1	6.4	2.2	4.1
Luxury	11.6	(1.2)	11.6	(0.4)	10.1	10.7	5.1	5.7
Pickup	9.9	(2.2)	10.6	(0.7)	9.5	9.1	5.3	6.1
SUV	8.2	(0.9)	7.2	(0.4)	7.4	6.9	3.4	4.3
Sporty	5.1	(2.4)	7.2	(0.8)	4.3	4.1	1.7	2.8
Van	9.6	(1.4)	8.6	(0.4)	8.3	8.7	4.3	5.2
Average	9.2	(0.6)	8.8	(0.2)	8.1	7.9	3.7	4.7

Note: The percentage price declines are at annual rates. Standard errors are in parentheses.

Table 10: Average Price Declines and Vintage Premia (in percent)

assembly plants, production during the first several weeks of the model year is typically below average as the plant “ramps up” output; further throughout the year, plants regularly close for holidays and in response to both supply and demand shocks. While these features can be easily incorporated into the model and would increase the volatility of production, they obscure the intuition of the model.

The model simulation can also be viewed in price-quantity space. In figure 18 we plot the time path of the quantity-price pairs over a single model year for a typical midsize car. The stars denote the weekly realizations from a single simulation from the model. The crosses denote the average observations from the data for the midsize sector. These quantity-price pairs are interpolated from monthly observations to obtain weekly points. As the figure illustrates, prices fall over the model year while sales initially start small, grow, and then decrease. Any model with a stable supply or demand curve will be unable to match the price and sales patterns in the data; however, the current model successfully replicates the basic horseshoe pattern. For the first half of the model year, demand is relatively stable while the marginal revenue curve rotates counter-clockwise toward the demand curve as the firm accumulates inventories and reduces the revenue tax. The shadow value of inventories (i.e. the supply curve) also shifts right. This generates decreasing prices and increasing sales. After about 30 weeks, the demand curve starts to shift to the left and continues falling for the remainder of the model year. This causes both prices and sales to fall during the second half of the product cycle.

To isolate the role inventory policy has on the optimal price paths, we resolve the model removing




Figure 18: Weekly (quantity,price) pairs over the model year. The stars are realization from one simulation from the model. The blue stars denote this year's model. The green stars denote last year's model. The red crosses (this year's model) and purple crosses (last year's model) are the observations in the data.

the revenue tax (i.e. setting $\tau = 0$). This simulation can be interpreted as allowing the firm to engage in a build-to-order rather than a build-to-stock inventory policy. In figures 19-22 we plot the price, sales, production, and inventory paths under a build-to-order policy while still allowing for multiple vintages. To ease comparison, we set the scales of the figures consistent with those in figures 14-17. Under a build-to-order policy, the firm sets a downward sloping price path that is far less dramatic than the price path with the revenue tax. Prices fall only \$838 (from \$19,369 to \$18,531) or 4.3 percent (3.1 percent at an annual rate) over the product cycle. For the first 40 weeks the firm essentially builds to order, producing and selling roughly 4000 cars per week (corresponding to operating two 40 hours shifts). As the firm nears the 52nd week, the firm moderates the price declines to dampen sales and accumulates modest inventories to carry over into the next model year. The firm then sells last year's vintage for only fourteen weeks into the following model year with an average vintage premium of 4.3 percent. Hence a model with just declining demand yields within-product-cycle price declines and vintage premia that are half of the declines generated by a model with a build-to-stock inventory motive. Furthermore, under the build-to-order policy neither sales nor inventories have the pronounced hump-shaped pattern seen in the data.




Figure 19: Monthly Prices.




Figure 20: Monthly Sales.




Figure 21: Monthly Production.




Figure 22: Monthly Inventories.

Simulated Prices, Sales, Production, and Inventories for a Typical Midsize Car By Model Year: $\tau=0$.

These patterns hold across the other market segments as well. As can be seen in table 10, under the build-to-order policy the average price decline is only 3.7. This is less than half of the average price decline implied by the build-to-stock models (8.1) or seen in the data (9.2). Furthermore the implied vintage premia with $\tau = 0$ is only 60 percent on average of the vintage premia with $\tau = 1$ (4.7 versus 7.9). While it is certainly the case that the demand for a particular vintage of new vehicle falls over the model year, this alone explains about half of the price decline in automobiles and misses salient comovements among prices, sales and inventories.

3 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 a market equilibrium model for a single vehicle line. While falling demand over the product cycle explains about one half these price declines, our model suggests that the other half of these declines is driven by build-to-stock inventory management.

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

Finally we view this paper as part of a larger research agenda to understand how firms adjust prices and production when faced with demand shocks. While the analysis presented above is entirely deterministic, in future work we plan to add persistent demand shocks and a richer cost structure to the model. We can then ask whether a negative demand shock will result in lower prices (e.g. rebates or other financial incentives) or in shutting down the plant. Given the important non-convexities in production scheduling at automobile assembly, it may be optimal at times to cut prices to keep the plants running while at other times it may be optimal to shut plants down for an “inventory adjustment”. Ultimately, we want to understand the shape of the supply curve for new vehicles and how it varies over the model year.

References

- [1] Ayres, I., Siegelman P., 1995. Race and gender discrimination in bargaining for a new car. *American Economic Review* 85, 304-321.
- [2] Berry, S., 1994. Estimating discrete-choice models of product differentiation. *Rand Journal of Economics* 25, 242-261.
- [3] Berry, S., Levinsohn J., Pakes A., 1995. Automobile prices in market equilibrium. *Econometrica* 63, 841-890.
- [4] Berry, S., Levinsohn J., Pakes A., 1999. Voluntary export restraints on automobiles: Evaluating a trade policy. *American Economic Review* 89, 400-430.
- [5] Bresnahan, T., Reiss, P., 1985. Dealer and manufacturer margins. *Rand Journal of Economics* 16, 253-268.
- [6] Bresnahan, T., Ramey, V., 1994. Output fluctuations at the plant level. *Quarterly Journal of Economics* 109, 593-624.
- [7] Chan, H., Hall, G., Rust, J., 2005. Price discrimination in the steel market. manuscript, University of Maryland.
- [8] 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.
- [9] Federgruen, A., Heching A., 1999. Combined pricing and inventory control under uncertainty. *Operations Research* 47, 454-475.
- [10] 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.
- [11] Goldberg, P., 1995. Product differentiation and oligopoly in international markets: The case of the U.S. automobile industry. *Econometrica* 63, 891-951.
- [12] Goldberg, P., 1996. Dealer price discrimination in new car purchases: Evidence from the consumer expenditure survey. *Journal of Political Economy* 104, 622-654.

- [13] 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.
- [14] 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.
- [15] Lazear, E., 1986. Retail pricing and clearance sales. *American Economic Review* 76, 14-31.
- [16] Nevo, A., 2000. A practitioner's guide to estimation of random coefficients logit models of demand. *Journal of Economics & Management Strategy* 9, 513-548.
- [17] Pashigian, P., 1988. Demand uncertainty and sales: A study of fashion and markdown pricing. *American Economic Review* 78, 936-953.
- [18] 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.
- [19] Pesendorfer, W., 1995. Design innovation and fashion cycles. *American Economic Review* 85, 771-792.
- [20] Petrin, A., 2002. Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy* 110, 705-729.
- [21] Ramey, V., Vine D., 2004. Tracking the source of the decline in GDP volatility: An analysis of the automobile industry. NBER working paper 10384.
- [22] Stokey, N., 1979. Intertemporal price discrimination. *Quarterly Journal of Economics* 93, 355-371.
- [23] Whiten, T., 1955. Inventory control and price theory. *Management Science* 2, 61-80.
- [24] Zettelmeyer, F., Scott Morton F., Silva-Risso J., 2001. Cowboys or cowards: Why are internet car prices lower? manuscript, Yale School of Management.
- [25] 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.