

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 Hall[§]

Yale University

February 14, 2005

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, the retail price (net of rebates and financing incentives) falls 9.4 percent over the 17-month model year. Concurrently, both sales and inventories are hump-shaped. To explain these time-series, we formulate a market equilibrium model for new automobiles. On the demand side, we use our micro-level data to estimate aggregate demand curves for each vehicle and vintage over the model year. On the supply side, we solve a dynamic programming model of an automaker who, while only able to produce one vintage of a product at a time, may accumulate inventories and sell multiple vintages of the same product simultaneously. Declining prices and hump-shaped sales and inventories are consistent with optimal inventory management. Further, 3/4 of the price decline is driven by inventory control.

Keywords: dynamic pricing, discrete-choice demand estimation, revenue management

JEL classification:

*We would like to thank Ana Aizcorbe, Steve Berry, Andrew Cohen, Gautam Gowrisankaran, Amil Petrin, John Rust, and John Stevens for helpful discussion and comments as well as seminar participants at numerous conferences and seminars. 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 other members of the staff of the Federal Reserve System.

[†]Division of Research and Statistics, Board of Governors, Mail Stop 82, C St. and 20th St. NW, Washington DC, 20551, e-mail: adam.m.copeland@frb.gov

[‡]Division of Research and Statistics, Board of Governors, Mail Stop 82, C St. and 20th St. NW, Washington DC, 20551, e-mail: wendy.e.dunn@frb.gov

[§]Department of Economics, Yale University, 28 Hillhouse Avenue, P.O. Box 208268 New Haven CT 06520-8268, phone: (203) 432-3566, fax: (203) 432-5779, e-mail: george.hall@yale.edu

Like most producers of consumer durable goods, automakers maintain inventories of their output and introduce new vintages of their products at an annual frequency. With the use of inventories, the introduction of a new vintage results in a period of time when two vintages of a product are sold simultaneously. This within-new-market competition is in addition to and potentially more important than the outside competition from used vehicles.¹ New goods of an older vintage are usually quite similar to those of the newer vintage and do not suffer from the asymmetric information problems inherent with the used goods market. Thus an important dimension to the automakers' pricing and production decisions is the annual introduction of new vintages of their products and the resulting period of time when multiple vintages of the good directly compete with one another in the new-good market.

This paper analyzes the pricing and production decisions of an automaker over the model year, taking into account the annual introduction of new vintages. We first document the within model-year pricing and production behavior in the new U.S. automobile market by vintage. 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 lifecycle of each vintage but also the competition across vintages in the new-vehicle market. The model's optimal decision rules are consistent with the pricing and sales patterns observed in the data.

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 since they record not only the actual transaction price (not the list price or invoice) 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, retail prices (net of rebates and incentives) fall 9.4 percent over the model year.
2. While 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 on average for 8.8 percent less than the new vintage.

¹Rust (1985) is a classic paper that provides a theoretical analysis of the interaction between new and used markets. Esteban and Shum (2004) provide an empirical analysis of the effect of secondary markets on automakers' behavior in equilibrium.

4. Both sales and inventories are hump-shaped over the model year.
5. All other things constant, higher inventories are associated with lower retail prices.

The fact that automobile prices decline within the model year has been documented previously by Pashigian, Bowen, and Gould (1995).² The second and third facts document well known characteristics of the automobile industry. For almost 5 months out of the year, automakers simultaneously sell two vintages of a model. Further, despite these vintages often being quite similar, the older vintage typically sells at a 8.8 percent discount. Our evidence suggests that the vintage premia we observe are driven by the depreciation of a model-year in the used-vehicle market. Using data on used-vehicle sales, we find that the difference in prices between two consecutive model years in the used market holding all other things constant is about 9 percent—consistent with vintage premia we see in the new car market.

The last two facts are not surprising, but in combination with declining prices over the model-year suggest a model with either stable demand or stable supply will be unable to replicate these 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 suggests that rightward shifts in the supply curve dominate changes in the demand curve early in the model-year. During the last 12 months the vehicle is sold, both prices and sales fall, suggesting leftward shifts to 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 revenue management with the economics literature on discrete choice models of product differentiation. On the demand side we employ the econometric methodology developed by the discrete choice literature (e.g. McFadden, 1973, Bresnahan, 1987, Berry, 1994, and Berry, Levinsohn, and Pakes, 1995, to name a few) to estimate consumers' preferences over automobiles. Our approach differs from this standard approach in three main ways. First, we have a better measure of prices as we use transactions prices instead of the usual list prices. Second, we estimate our demand-side model at the quarterly frequency rather than annual; thus we estimate how the demand curve shifts throughout the model year. Third, we allow consumers to choose among multiple vintages of the same model as well as across models. Using our estimates of consumers' preferences, we compute average demand curves for each automobile market segment (e.g. compact cars) and vintage over the automobile product cycle.

A main result from this analysis is that demand curves for most vehicles shift significantly leftwards over the second half of the model year, with small changes to their slope. Further, we estimate that cross-

²More recently, Corrado, Dunn and Otoo (2004) construct an aggregate price index for light motor vehicles that also shows a decline in price over the model year.

price elasticities between models of different vintages are quite small. Hence, despite the similarities between vintages of the same model, consumers view these products as poor substitutes. 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.

Taking these demand curves as given, we formulate and solve a dynamic structural model with overlapping vintages in which the 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 only produces 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 how many units of the current vintage to produce and the optimal prices for the vintages in stock.

A significant aspect of the automotive market 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 observing the vehicle they are considering purchasing as well as taking possession of their vehicle without delay. Consequently, part of the automaker's problem is ensuring there are vehicles on dealer lots across the national market.³ We do not model the interaction between the manufacturer and the automotive dealer, but rather assume they are the same firm. We incorporate the distributional cost associated with dealerships into the firm's problem in ad hoc fashion, 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, raising the tax the firm pays per transaction.

After calibrating the model, we are able to replicate the decline in prices over the model year along with hump-shaped sales and inventories, as seen in the data. Early in the model year, the automaker sets the vehicle price high to dampen sales and so accumulates 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 (i.e. it lowers the revenue tax). Over the remainder of the model year, our estimates of leftward shifting demand lowers the shadow value of inventories, resulting in a 12.5 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 resolve the model removing the revenue tax, thus allowing the firm to manage inventories

³The automotive trade press often mention the necessity for a showroom presence when discussing manufacturers' inventories. In 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-day supply [i.e. inventories over sales] because they want to display all the models".

on a “build-to-order” basis. In this case, retail prices fall over the product cycle but by only one-fourth of the magnitude we see in the data. Hence, unlike previous work that attributes falling prices to fashion or demand uncertainty, our theory implies that about three-fourths of the magnitude of the price declines are the result of optimal build-to-stock inventory management. With overlapping vintages, falling demand curves alone over the product cycle will not generate falling prices concurrently with hump-shaped sales and inventory. In order to match these facts, there must be shifts over the product cycle in both supply and demand. We have not proven that our theory is the only theory that can explain facts, but we can show that falling demand alone cannot explain both of these facts.

1 Data

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, and provides information on prices by model and model-year as well as the distribution of sales across model-years by model. The second data source contains information on total sales in North America by country and model, as well as 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 United States, 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 United States 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 (e.g. 2000 Honda Accord) sold in the United States, and covers the period from January 1999 to January 2004. Among other variables, the dataset includes information on the number of transactions recorded in the JDPA dataset, the average transaction price, the average cash rebate, and details of the average financial package customers received. JDPA attempts to precisely measure the transaction price of a vehicle. This measure includes the price of accessories (e.g. roof racks) and transportation costs, and excludes aftermarket options, taxes, title fees, and other documentary preparation costs. Further, JDPA

adjusts this price to account for instances when a dealership under- or over-values 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 capture loans customers received from all financial institutions, conditional on the loans being 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 accomplish this by first computing the net present value of the average amount financed given the interest rate and loan term in the data. We then compute the net present value of financing the same average amount using the average interest rate reported for 48 month new car loans at commercial banks.⁴ The value of the manufacturer's financial incentive is then defined as the difference between the net present value of the amount financed given the terms in the data versus those at a commercial bank. Finally, we convert the market price into 2000 dollars using the BEA's personal consumption deflator.

In addition to providing the real market price of vehicles by model and model-year, the data from JDPA also provide us with the distribution of sales by model-year for every model. Using the monthly data they provided on the number of transactions by model and model-year, we compute the fraction for which each available model-year accounts of a model's sales.

We linked the JDPA data to information on General Motors, Ford, and Chrysler's U.S. sales and North American production, which we obtained from Ward's Communications. We excluded foreign manufacturers as it is difficult to measure overseas production. The sales data on 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 based on the monthly model-year distributions in the JDPA sample. Using information from Ward's on model change-over 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. The above

⁴These data are published by the Board of Governors in their G.19 Consumer Credit release

work results in a dataset with monthly observation 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. by model-year from 1999 to 2003.

1.2 Empirical Observations

As stated earlier, by examining these data we are able to observe several stylized facts that hold across models and model years. To provide illustrative examples, figures 1-8 show plots of the sales, price, production and inventory data for a midsize car and a pickup truck. The steady decrease in price over the sales cycle is immediately evident for both vehicles shown in the charts. In the 2000 model year, the average market price for the midsize car falls by 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 year. 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. We refer to this difference in price as the “new vintage premium.” The size of this premium varies, with the midsize car averaging almost 7 percent while the premium for the pickup averages about 9 percent. Turning to the figures on sales and inventories, both the midsize car and the pickup illustrate a hump-shaped profile of sales and inventories.

To observe the first stylized fact, the within-year price declines, more generally, figure 9 illustrates the aggregate matched-model price indexes for successive model years as constructed by Corrado, Dunn, and Otoo (2004). This price index was constructed using 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 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. There is little variation in the number of months sold across vehicles; the mean length of the automobile product cycle has a standard error of only 0.02 .

Turning to the third stylized fact, table 1 reports the average new vintage premium by market segment and model year, weighted by sales. Looking at Midsize cars, the table shows that this category accounts for 20 percent of all vehicles sold by the Big Three over this time period. Further, the table illustrates that the new vintage premium of 2000 model-year Midsize cars over their 1999 model-year counterparts is 10.0




Figure 1: Average transaction prices (net of rebates and financing incentives) for a new midsize car.

Source: J.D. Power and Associates and authors' calculations

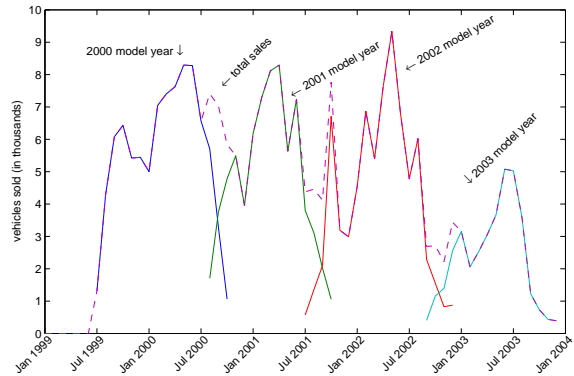


Figure 2: Monthly sales of a new midsize car. The dashed line is the sum of sales across model years.

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




Figure 3: Monthly production of a midsize car.

Source: Wards Communications and authors' calculations




Figure 4: Monthly inventories of a midsize car.

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




Figure 5: Average transaction prices (net of rebates and financing incentives) for a new pickup by model year.

Source: J.D. Power and Associates and authors' calculations




Figure 6: Monthly sales of a new pickup by model year. The dashed line is the sum of sales across model years.

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




Figure 7: Monthly production of a pickup by model year.

Source: Wards Communications and authors' calculations




Figure 8: Monthly inventories of a new pickup by model year.

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




Figure 9: Matched-model Price Indexes by Model Year.

percent. In the sample, the new vintage premium averages 8.5 percent for Midsize cars with a standard error of 0.4.

Though it varies quite a bit across both market segments and time, overall the new vintage premium averages 8.8 percent in our sample. This premium is highest for Luxury cars and Pickup trucks, and lowest for Compact cars, SUVs 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 a large 13 percent. Not surprisingly, over the time period in our sample the Big Three placed the largest incentives on 2003 model-year vehicles, and these incentives were disproportionately offered in the latter half of the product cycle.

Of course, one might argue that the new vintage premium simply reflects improvements in quality or additional features of a vehicle. For example, the 20 percent vintage premium recorded for 2004 model-year Pickup trucks reflects, in part, a quality improvement made to Ford's F-series pickup truck.⁵ 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 these vehicles the downward sloping price pattern is still apparent. To further investigate this fact, we looked at the new vintage premium for the subsample of vehicles

⁵Alternatively, with the introduction of a cheaper base model, the vehicle premium could be biased downwards (e.g. -7.8 percent for 2004 Sporty cars).

Market Segment	% of Sample	Model Year										All	
		2000		2001		2002		2003		2004			
Compact	6	5.9	(1.0)	6.9	(0.8)	6.7	(0.7)	6.9	(1.1)	11.0	(1.1)	7.1	(0.5)
Midsize	20	10.0	(0.9)	5.8	(0.9)	6.1	(0.4)	7.5	(0.5)	11.9	(1.0)	8.5	(0.4)
Fullsize	6	9.9	(1.6)	6.2	(0.8)	7.8	(1.2)	8.2	(1.2)	9.2	(1.2)	8.3	(0.6)
Luxury	15	11.0	(0.8)	11.6	(0.6)	9.8	(0.6)	14.1	(0.7)	11.1	(1.3)	11.6	(0.4)
Pickup	8	10.7	(1.4)	9.8	(0.9)	6.3	(1.2)	8.6	(1.1)	20.0	(3.0)	10.6	(0.7)
SUV	25	5.4	(0.9)	0.4	(1.0)	10.1	(0.8)	8.8	(0.6)	10.9	(0.6)	7.2	(0.4)
Sporty	7	2.6	(1.1)	7.5	(1.2)	3.3	(1.0)	28.9	(4.5)	-7.8	(2.9)	7.2	(0.8)
Van	13	7.7	(0.5)	11.8	(0.8)	3.6	(0.7)	9.1	(0.9)	12.4	(0.9)	8.6	(0.4)
All	100	8.6	(0.4)	7.2	(0.3)	7.1	(0.3)	9.1	(0.3)	13.1	(0.5)	8.8	(0.2)

Note: Standard errors are in parenthesis

Table 1: The Average ‘New Vintage Premium’ by Market Segment and Model Year

excluding those that experienced a major re-design. We determined when a model received a major re-design by using data from Wards 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 re-design. We re-computed the new vintage premium for these vehicles and found that the average premium is not much different than the figures reported in table 1.

The fourth stylized fact states that both sales and inventories for a particular model and model-year exhibit a humped-shaped pattern. Figure 10 plots the Big Three’s aggregate sales by model year, and shows the distinctive hump-shaped that sales follow over the product cycle. The contour of aggregate sales, however, confound the evolution of sales over the product cycle with calendar effects, because not all vehicles are 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 for how many months the model-year has been sold. If $1_{t=1}$ is equal to one, then at this date the associated vehicle is in its first month of sales. For model-years that are sold for more than 14 months, we define a dummy variable 1_{15+} which is equal to one for sales taking place more than 14 months from 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. Table 2 reports the estimated coefficients from running these two regressions.




Figure 10: Big Three Sales By Model Year

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 2: The Shape of Sales and Inventories by Model and Model-Year Over the Product Cycle




Figure 11: Days Supply Over the Product Cycle.

Looking first at the regression using sales, the estimated coefficients imply that sales do indeed follow a humped-shaped path. After rapid sales growth in the first three months of sales, the level of sales then slowly climbs up to its peak six to eight months after the vehicle’s debut, before declining. This same pattern can be seen with 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. Figure 11 charts this ratio for three market segments over the product cycle. 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 months the firm could continue to sell cars only using the their stock of inventories coming into the month, assuming future sales are equal to the current month’s flow.

As illustrated in the figure, the interesting pattern in this measure of days supply is the spike in the first month, followed by a decline over the product cycle. The spike is a consequence of automakers having small sales at the beginning of the product cycle. The decline in days supply over the remainder of the product cycle signifies that firms maintain an increasingly small stock of inventories relative to the pace of sales.

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 price 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 above provide a measure of the deviations from the usual contour of inventories over the product cycle. We measure if 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 CPI data to show that at the aggregate level prices for new cars decline between December and the following September of the model year. While the magnitude of the within model-year price variation declines have fallen between 1954 and 1989, Pashigian, Bowen, and Gould find these price declines are larger for luxury and speciality cars than for compact 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.⁶ In contrast, we find only marginal evidence of a particular pattern in the vintage premium across market segments. As shown in table 1 while luxury cars command the highest vintage premium, other fashion-oriented vehicle types such as SUV's and Sporty automobiles do not command premiums above the more plain Compact and Midsize automobiles.

As such, we posit that within model-year price declines are driven more by the used-vehicle market than by fashion. As an example, consider the case of a 2000 model-year vehicle produced at the end of the product cycle and a 2001 model-year produced at the beginning of the product cycle. Although these two vehicles may have produced just a few weeks apart, we expect that the 2000 vintage will be perceived on the used market to have been on the road 12 months longer than the 2001 vehicle.

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

⁶See Lazear (1986) and Pashigan (1988) for more detail on the optimal pricing of fashion goods.

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

Table 3: Age and Odometer Coefficients from the Used Vehicle Price Regression

one.⁷ Table 3 shows the resulting coefficients on age and odometer reading, both of which are statistically significant with greater than 99 percent confidence. 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, however, the coefficient on age implies that, even after controlling for the odometer reading and other vehicle characteristics, a higher model age, (i.e. 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 year decreases the value of a used vehicle by 9.3 percent, 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 a vehicle’s value of the used 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 above. On the demand-side we draw upon the existing discrete-choice literature. We assume households solve a static utility maximization problem each quarter. We then estimate the parameters that household’s optimization problem. Using these point estimates we compute the own-price and cross-price elasticities by model and model-year over the product cycle.

On the supply-side we assume the automaker takes as given the estimated demand curves and solves a dynamic profit maximization problem. As the automaker is able to hold inventories, there are points in time where the automaker is able to sell two vehicles: last year’s vintage and this year’s vintage. While we formally estimate the model parameters on the demand side, the model parameters on the supply side are chosen to match the key features of the firm’s cost structure. We derive decision rules governing the production and pricing of vehicles over the model-year. Via numerical simulations, we demonstrate that the empirical regularities documented above are consistent with our derived decision rules.

⁷Because 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 less. This reduces the variation in price across vintages of the same model due to changes in unobserved characteristics.

2.1 Automobile Demand

2.1.1 Overview

The demand for automobiles is modelled using a discrete-choice framework. Following Berry, Levinsohn, and Pakes (1995), henceforth BLP, we construct the demand system by aggregating over the discrete choices of heterogenous individuals. The utility derived from choosing an automobile depends upon the interaction between the consumer's characteristics and the product's characteristics. Consumers are heterogenous 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) to be denoted by X , and those that are unobserved by the econometrician (such as styling or prestige), to be denoted by ξ . Drawing from the nested logit literature, we also incorporate a correlation in consumer's tastes for vehicles of the same model year. We divide vehicles into $G + 1$ mutually exclusive groups (i.e. model years), $g = 0, 1, 2, \dots, G$, where the outside good is the sole member of group 0. We specify the indirect utility derived from consumer i purchasing product j , dropping the time subscript, as

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

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 independent and identically distributed standard normal. The parameter σ_k captures the variance in consumer tastes. The term α_i measures a consumer's distaste for price increases. As in Berry, Levinsohn, and Pakes (1999), we assume that $\alpha_i = \frac{\alpha}{y_i}$, where y_i is a draw from the income distribution and α is a parameter to be estimated. We assume the distribution of household income is lognormal and, for each year in our sample, estimate its mean and variance from the Current Population Survey (CPS). The second-to-last term in equation 1 captures correlations in 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 identically and independently distributed extreme value.

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

2.1.2 Implementation

Using the data described in the previous section, we aggregate sales and prices to the quarterly level, and so observe 20 markets. We do not use the data at a monthly frequency because there is a significant amount of volatility in monthly sales due, in part, to intertemporal substitution. As such, BLP's static utility maximization approach is better suited for 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.

The vehicle characteristics that we use include a set of model dummies, a measure of acceleration (horsepower divided by weight), vehicle dimensions, a measure of safety and fuel efficiency.⁸ As done in previous research, quantity sold and the transaction price are linked 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 observable characteristic vector 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, through ζ we also incorporate 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 a measure of market size for the year. We assume that one quarter 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

⁸Information on vehicle characteristics was taken from various years of the Automotive News's Market Data Book

⁹We modified the programs provided in Nevo (2000) to estimate the demand system. An important 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.

disturbance ξ_j , using the assumption that ξ is uncorrelated with the vehicle characteristics X , or

$$E[\xi(\theta)|X] = 0. \quad (2)$$

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

$$\delta_j = X_j\beta - \xi_j + \rho \ln(s_{j/g}) \quad (3)$$

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 both with price and $s_{j/g}$, there is an endogeneity problem. 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 diminished effect. These instruments are based on competing products' characteristics, which change at the model-year frequency, while our price and quantity data vary at the quarterly level. As such, we augment the set of instruments to include the level of beginning-of-period inventories by model.

Inventories are valid instruments in a dynamic setting as theory suggests they are strongly correlated with both price and market share while, under some mild assumptions, remaining uncorrelated with ξ . Inventories serve as good instruments because they are a function of supply shocks such as labor strikes, regional power outages, and shortages of parts. In a dynamic setting, however, inventories may also be correlated with demand shocks. This is not the case, though, if our instruments are beginning-of-period inventories and we assume that a vehicle's unobserved component is composed of a constant and an i.i.d. term. The constant term, like the observable characteristics of the vehicle, is known to the firm, while the i.i.d. component of ξ is only observed after the pricing and production decisions have been made. In this setting, the beginning-of-period inventories are a function of last period's ξ , but not this period's ξ . In contrast, if the firm learns the value of ξ over time or ξ is serially correlated, beginning-of-period inventories would not be appropriate instruments because they would be correlated with both last and this period's ξ .

2.1.3 Demand Results

In this section we display the results from estimating the demand side of our model. We then discuss the own and cross price elasticities implied by our parameter estimates.

Characteristic	Mean	Std. Dev.	Min	Max
Price (\$)	27,651	12,594	9,176	87,552
Horsepower/weight (hp/lbs)	0.053	0.015	0.029	0.149
Height (in.)	62.5	8.9	47.0	81.6
Size (in. sq.)	14,229	1,624	9,852	18,136
Airbag	0.29	0.45	0	1
Miles/dollar	0.125	0.028	0.056	0.287

Table 4: Average (Sales-Weighted) Characteristics

Note: Airbag is a dummy variable equal to one if driver, passenger, and side airbags are available

In addition to model dummies, the product characteristics we use in our demand analysis are price, horsepower divided by weight (a measure of acceleration), height, size (length times width), a measure of safety (whether driver, passenger and side airbags are available), miles per dollar (miles per gallon divided by real price of gas per gallon), and the number of quarters for which a vehicle has been sold. Table 4 reports various statistics for the product characteristics in our data sample. For our current demand analysis, we used data covering General Motors, Ford and Chrysler from January 1999 to January 2004, which include 638 unique model and model-year vehicles. As shown in the table, there is much variation in these five characteristics. In our sample, the real price of vehicles ranged from a little over \$9,000 to almost \$90,000 and the size of the largest vehicle is almost double that of the smallest. In our sample, 29 percent of vehicles offer driver, passenger and side airbags, though these instances are concentrated at the end of the our sample period.

The relevant parameter estimates are presented in table 5. Given their large number we do not report our fixed effects estimates. Rather, we report those estimates that measure the heterogeneity in consumer's tastes (σ) along with the estimates of the substitutability of vehicles across model-years and of consumer's distaste for price. The constant term, airbag (measure of safety), and intro (number of quarters since introduction) are not statistically significant. We estimate that consumers are quite heterogenous, however, in their taste for horsepower over weight (measure of acceleration), height, size, and the miles per dollar (measure of fuel efficiency).

Both the model-year nest and price coefficient 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. Both the estimate of ρ and the magnitude of the price coefficient are more easily interpreted by examining the implied own-price and cross-price elasticities. Table 6 reports the average of own-price elasticities of individual vehicles averaged across market segments, quarters, and

Parameters Std. Deviations (σ)	Variables	Parameter Estimate	Standard Error
	Constant	0.43	0.271
	HP/Weight	4.65	1.360
	Height	1.94	0.454
	Size	0.63	0.188
	Miles/dollar	10.03	4.137
	Airbag	0.89	0.64
	Intro	-0.12	0.091
Model-Year Nest	ρ	0.61	0.013
Distaste for Price	α	36.89	0.292

Table 5: Parameter Estimates

vintages, where the vintage label signifies whether the vehicle is the newest model-year available or not.

The quarters in table 6 do not correspond to the definition of a calendar quarter. Rather, 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 using this new grouping of months.

The own-price elasticities generated by our parameter estimates range between 6 and 12, implying that manufacturers' face quite elastic demand. Looking at the first quarter a car is sold, our results state that a 1 percent price increase for a typical compact car (roughly \$140) causes an 8.5 percent fall in sales, holding everything else equal. Our measure of elasticities fluctuates across quarters, though not systematically. 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. Verboven (1996), however, reports elasticities inline with our results. In addition, it is not surprising that our elasticity estimates are higher, because previous research estimated own-price elasticities among models, which is a higher level of aggregation compared to our data. Indeed, when we re-estimate the parameters using data aggregated to the model level, the implied own-price elasticities fall within 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. Table 7 reports our estimates of the cross-price elasticities between two vintages of the same model averaged over quarters and across market segments.

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 quarter (i.e. August through January). However, a fair number of vehicles are introduced at other times in the year, allowing us to compute cross-price elasticities through-

Vintage	Market Segment	1st Quarter	2nd Quarter	3rd Quarter	4th Quarter
New	Compact	8.5	7.9	8.0	9.1
	Full	11.9	10.3	9.8	11.6
	Luxury	9.7	9.6	9.1	10.5
	Midsize	10.7	10.5	9.6	10.5
	Pickup	11.1	11.4	9.2	10.3
	SUV	10.0	9.6	9.0	9.1
	Sporty	10.2	9.9	9.0	9.2
	Van	10.2	10.4	9.3	10.2
Old	Compact	8.1	7.2	9.5	9.5
	Full	10.7	10.6	9.8	9.5
	Luxury	9.6	9.1	9.4	8.9
	Midsize	10.5	10.7	11.2	11.7
	Pickup	11.2	11.0	10.4	10.2
	SUV	9.6	9.4	8.9	12.0
	Sporty	10.2	10.0	10.0	9.4
	Van	10.0	9.4	9.1	5.7

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

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	0.01	0.01	0.01
	Luxury	0.02	0.01	0.00	0.01
	Midsize	0.02	0.00	0.00	0.03
	Pickup	0.02	0.01	0.00	0.01
	SUV	0.09	0.02	0.02	0.10
	Sporty	0.03	0.01	0.02	0.03
	Van	0.02	0.01	0.00	0.00
Old to New	Compact	0.01	0.03	0.03	0.01
	Full	0.01	0.03	0.04	0.01
	Luxury	0.01	0.03	0.01	0.00
	Midsize	0.01	0.02	0.01	0.01
	Pickup	0.01	0.03	0.01	0.02
	SUV	0.05	0.12	0.02	0.00
	Sporty	0.02	0.04	0.02	0.04
	Van	0.01	0.02	0.02	0.01

Table 7: Cross-Price Elasticities Between Vintages of the Same Model by Market Segment, Quarter

out the year. The upper portion of table 7 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. Our estimates of the cross-price elasticities are more than an order of magnitude smaller than the own-price elasticities, implying 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 utility function. We rejected a specification of utility where groups are defined at the model level (i.e. 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 small cross-price estimated elasticities are consistent across several different utility specifications.¹²

2.2 Automobile Supply

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 sub-firm or profit center. Hence, an automaker is modelled as a collection of dynamic programs that can be solved independently from each other. Second, we integrate the dealership into the automaker and consider an unified pricing decision. Third, we abstract from issues of bargaining and price discrimination by assuming 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 side and marketing side of the business. But since these issues are not central to understanding the facts presented above, we defer further consideration to other papers.¹³

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). Like most of the papers in this literature we assume that the good must be sold by a fixed deadline; but we extend this literature to allow the firm to sell two vintages simultaneously and to have a cost structure of producing vehicles with several nontrivial

¹⁰Ana Aizcorbe suggested there may be geographical reasons behind 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 may substitute between vintages may be limited.

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 with a greater number of parameters than the specification described in the paper (e.g. random coefficients on a set of model dummies). These alternative specifications demanded too much of the data.

¹³For 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 Risso (2001).

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

2.2.1 Supply Model

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 starts 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 since this year's vintage becomes last year's vintage at the end of the T^{th} week. Each week the firm must decide: 1) how many vehicles of the current model year to produce, q_t ; 2) the number days the plant will operate, D_t , the number of shifts 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 weekly sales, s_t^j , for each of the two vintages depend on their own price and the price of the other vintage:

$$s_t^j = \mu_t^j - \eta_t^j \log(p_t^j) + \phi_t^{ji} \log(p_t^i) \quad \text{for } j, i = \{\text{this, last}\} \text{ and } i \neq j. \quad (4)$$

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 inventoried at the end of period t carried over into period $t + 1$. Current production is not available for immediate sale, so sales can only be made from the beginning-of-period inventories:

$$s_t^j \leq I_t^j. \quad (5)$$

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

$$I_{t+1}^{this} = I_t^{this} + q_t - s_t^{this}. \quad (6)$$

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

$$I_{t+1}^{last} = I_t^{last} - s_t^{last}. \quad (7)$$

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:

$$I_1^{last} = I_T^{this} + q_T - s_T^{this}. \quad (8)$$

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, 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 of length h in hours. 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:

$$q_t = D_t \times S_t \times h_t \times LS. \quad (9)$$

While the production function is linear, the firm faces several important non-convexities due to its labor contract.¹⁴ The average straight-time, day-shift wage at these plants about is on average \$27 an hour plus benefits. Workers on the second (evening) shift receive a 5 percent premium. Any work in excess of eight hours in a day and all Saturday work is paid at a rate of time and an half. Employees working 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 the worker receives for the hours s/he actually worked. The firm can shut down the plant temporarily at no cost.

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. This implies the firm's week t cost function is:

$$\begin{aligned} c(D_t, S_t, h_t | q_t) = & q_t \times \gamma + (w_1 + I(S_t = 2)w_2) \times (D_t h_t n + \max[0, 0.85(40 - D_t h_t) n_t]) \\ & + \max[0, 0.5 D_t (h_t - 8) n_t] + \max[0, 0.5 (D_t - 5) 8 n_t] \end{aligned} \quad (10)$$

where w_1 and w_2 are the hourly wage rates paid to the first-shift and second-shift workers, respectively. So the first term with the bracket represents the straight time wages paid to the production workers. The second, third, and fourth terms capture the 85 percent rule for short-weeks and the required overtime premium. The first term outside the bracket is the per-vehicle cost of building a vehicle; this term incorporates

¹⁴For further discussion of the institutional details of the labor contracts in automobile manufacturing, see Bresnahan and Ramey (1994), Hall (2000), or Ramey and Vine (2004). Hall and Ramey and Vine assume sales follow an exogenously specified Markov process and thus the firm solves just a cost minimization problem. In this paper, we solve a profit maximization problem in which the firm must decide both the price it wishes to charge and the number of vehicles it wishes to produce.

all costs (e.g. materials, energy, transaction) that do not depend on the allocation of production over the week. Let $D_t = 0$ if and only if $S_t = 0$. Thus, the cost function is piecewise linear with kinks at 1 shift running 40 hours per week and 2 shifts running 40 hours per week. This implies that the firm will minimize average costs by operating the plant at either one or two shift for 40 hour 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

$$E \left\{ \sum_{t=1}^T \left(\frac{1}{1+r} \right)^{t-1} \left\{ p_t^{last} s_t^{last} (1 - (s_t^{last}/i_t^{last})^2) + p_t^{this} s_t^{this} (1 - (s_t^{this}/i_t^{this})^2) - c(D_t, S_t, h_t | q_t) \right\} + \left(\frac{1}{1+r} \right)^T V(I_1^{last}, 0, 1) \right\} \quad (11)$$

subject to (4)-(9) and where $c(D, S, h | q)$ is given by (10). The terms $(s_t^{last}/i_t^{last})^2$ and $(s_t^{this}/i_t^{this})^2$ are "revenue taxes" the automaker must pay if the sales-to-inventory ratio is large. This term captures the distributional costs that automakes faces, as described previously in the introduction. When inventories are low, it is harder for potential customers to observe and gauge the vehicle (e.g. to test-drive it), making it more costly to match the buyer to a vehicle and consummate a sale. The tax effectively disappears when the sales-to-inventory ratio is small. 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 for the firm at week t that holds I^{last} of last year's vintage and I^{this} of this year's vintage in inventory. Then the firm's value function for this vehicle can be written:

$$V(I^{last}, I^{this}, t) = \max_{p^{this}, p^{last}, q} \left\{ p^{last} E(s^{last}) (1 - (E(s^{last})/i^{last})^2) + p^{this} E(s^{this}) (1 - (E(s^{this})/i^{this})^2) - \min_{D, S, h} c(D, S, h | q) + \frac{1}{1+r} EV(I^{last} - s^{last}, I^{this} + q - s^{this}, t+1) \right\} \quad \text{for } t = 1, \dots, T-1 \quad (12)$$

subject to (4), (5), and (9) with $c(D, S, h | q)$ given by (10).

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

$$V(I^{last}, I^{this}, T) = \max_{p^{this}, p^{last}, q} \left\{ p^{last} E(s^{last}) (1 - (E(s^{last})/i^{last})^2) + p^{this} E(s^{this}) (1 - (E(s^{this})/i^{this})^2) - \min_{D, S, h} c(D, S, h | q) + \frac{1}{1+r} EV(I^{this} + q - s^{this}, 0, 1) \right\}. \quad (13)$$

We solve the T Bellman equations numerically with a combination of backward recursions and policy iteration.

2.2.2 Supply Results

To illustrate the dynamic pricing and production behavior implied by our model, we solve a discrete approximation to (12) and (13) numerically for a set of parameter values chosen to match a typical car in the midsize segment. 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 in the model to holding inventories. The key specification is the law of motion for sales for the two vintages, equation (4). For these parameters we use the estimated own-price semi-elasticities and cross-price semi-elasticities implied by the demand estimation detailed in section 2.1. We set μ_t to match the average price and quantity of each market segment for each month. We then interpolated the demand curves to the weekly frequency.

To parameterize the cost function, we select values that are both sensible and that imply model's average prices and quantities sold over the product cycle are equal to the sample averages. While we would have preferred to econometrically estimate these parameters, computational issues make it infeasible. The line speed at most North American assembly plants that produce midsize cars are set between 40 and 60 cars per hour; thus we fix the line speed to 50 cars per hour. We set n to 1300 workers per shift shift, so the firm has 2600 workers employed. We read the wages off the union contract: $w_1 = \$27.00$ per hour and $w_2 = \$28.35$ per hour. Finally we set γ , the per vehicle cost equal to \$17,500.

Using these parameter values, we solve the dynamic program given by (12) and (13) via a combination of backward recursions and policy function iteration. Specifically we merge the T value functions into a single time-invariant Bellman equation:

$$V(I^{last}, 0, 1) = \max_{\{p_t^{this}, p_t^{last}, q_t, D_t, S_t, h_t\}} \left\{ E_1 \left(\sum_{t=1}^T \left(\frac{1}{1+r} \right)^{t-1} \left(p_t^{last} s_t^{last} (1 - (s^{last}/i^{last})^2) + p_t^{this} s_t^{this} (1 - (s^{this}/i^{this})^2) - c(D_t, S_t, h_t | q_t) \right) + \left(\frac{1}{1+r} \right)^T V(I_T^{this} + q_T - s_T^{this}, 0, 1) \right) \right\}.$$

To solve for the fixed point, we employ an algorithm suggested by John Rust: 1) Guess an initial value for $V(I^{last}, 0, 1)$; 2) Solve the T Bellman equations in (12) and (13) via backward recursions; 3) compute a new value for $V(I^{last}, 0, 1)$ via policy iteration; 4) repeat steps 2) and 3) until a fixed point is reached.

Due the non-convexities in the cost function, we solve for both the optimal level of output and the cost minimizing production schedule via grid search. We allow weekly production, q to take on values between 0 and $6 \times 10 \times 2 \times 50 = 6000$ 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 and 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 there is more curvature in the value function. For each of the 676 inventory pairs, we maximize the right hand side of equations (12) and (13) over each sales price and level of output. Expectations are approximated by Gaussian quadrature and points off the two inventory grids are approximated via bi-linear interpolation. We allow the two sales prices, p^{last} and p^{this} , to take on any positive value such that quantity demanded remains positive. Finally, we imposed only modest seasonality on production, assuming the plant closes for two weeks in July for a model changeover (weeks 51 and 52).

To illustrate the solution of the model, we plot in figure 12 the partial derivative of the value function for week 26 with respect to inventories of the current model year. (Other weeks are qualitatively similar.) This figure illustrates the marginal increase in the firm's net worth from an additional unit of inventory at each point in the state space for week 26. This is often referred to as the *shadow value of inventories*. The shadow value of inventories is a decreasing function of the level of inventories. When the the level of inventory for this year's model is close to zero, an additional vehicle of this year's model is worth \$18,750 to firm; however at the upper bound of the inventory grid, the shadow value of an additional vehicle from this model year is only worth a little over \$18,100 to the firm – roughly the average cost of producing a vehicle when the plant operates two 40 hours/week shifts. 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 \$17,500 provides a lower bound on the marginal cost.

Since 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, this optimal price for this year's model is:

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

where V_2 denotes the derivative of the value function with respect to the second argument. This is just 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 13. The pricing rule is almost to the shape of the shadow value. Holding all other things constant, the optimal price is a decreasing




Figure 12: Week 26 Shadow Value of Inventories for This Year's Vintage.




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

function of the level of inventory. This is our fifth fact. Further the price functions are consistent with the findings of Zettlemeyer, Scott-Morton, and Silva-Risso (2003) that the average retail price at a dealership with ample inventory is about \$230 per car less than the retail price at a dealership with low inventory.

In figure 14 we plot slices of the pricing rules for different weeks in the product cycle holding the inventory of the competing vintage fixed. For the current model year, the pricing rule are downward-sloping and shifting down in a parallel manner over time. Once a vehicle becomes last year's model, the firm's problem essentially becomes a finite horizon, no-replenishment pricing problem with independent demand over time. At any week, prices are a decreasing function of the level of inventory but the pricing curve shifts down over time. Since inventories are monotonically decreasing over time (since there is no replenishment) prices may go up or down depending on the evolution of inventories.

In figures 15-18 we plot a simulation from the model for five 52 week model years time-aggregated to a monthly frequency. Since the model is deterministic each of these simulations is identical, but these graphs are designed to be analogous to figures presented in section 1; however note that figures 1-4 are for a particular midsize while we use estimated demand curves in 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. Further the automaker wishes to build up inventories in order to reduce this tax in the future. Hence the automaker sets production high (6000 vehicles per week) and sets prices high early on in the year to dampen down sales and allow inventories to accumulate. Once




Figure 14: Optimal Price Rule Over the Product Cycle, Holding the Inventory level of the Competing Vintage Fixed.

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. As the model year progresses, demand for the vehicle starts to decrease as the demand curve shifts down further exacerbating the fall in prices. For twenty-eight weeks (a little over half a year) the automaker sells both vintages simultaneously and prices fall from \$19,853 to \$17,366 (a drop of \$2,487 or 12.5 percent) over the product cycle. This results in a new vintage premium of 7.9 percent, which is within two standard deviations of the 8.5 percent new vintage premium for a typical midsize car, as reported in table 1.

For most of the year, the automaker produces 4000 vehicles per week running two eight-hour shifts for five days per week. However, comparing figures 3 and 7 to figure 17, it is clear that the model has several counterfactual implications for the time-path of production. First, for most vehicles (including the two we plot in section 1) production during the first several weeks of the model year is typically below average as the plant “ramps up” output. In contrast, the model predicts that production is highest in the first weeks of model year. Adding a constraint restricting the quantity of production early in the model year would fix this counterfactual implication and improve the fit of the model by raising the shadow value of inventories early in the model year; nevertheless we report results without this ad hoc constraint because




Figure 15: Simulation of monthly retail prices over the model year.




Figure 16: Simulation of monthly sales over the model year.




Figure 17: Simulation of monthly production over the model year.




Figure 18: Simulation of monthly inventory holdings over the model year.

we feel this constraint would obscure the primary mechanism in the model. Second, simulated production is considerable smoother than production in the data. Adding i.i.d. shocks to the demand curves and imposing that the plant close for weeks corresponding to holidays (e.g. the week between Christmas and New Year’s Day) increases the volatility of production and the realism of the model; however we feel adding these features to the model reduces the clarity of the intuition behind the results.

In figure 19 we plot the time path of the quantity-price pairs over the model year in “demand and supply” space. The stars denote the weekly realizations from a single simulation from the model. The crosses denote the average observations for the Midsize sector observed in the data. The quantity-price from the data are interpolated from monthly observations to obtain weekly points. As the figure illustrates




Figure 19: 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.

prices fall over the model year while sales initially start small, grow, and then decrease over the model year. This implies that the supply curve must be shifting to the right early on in the model year, and the demand curve must be shifting to the left in the second half of the model year.

In our model, the automaker sets the vehicle price high early in the model year to dampen sales and accumulate inventories. Building up inventories reduces the cost of carrying out a transaction by lowering the revenue tax and shifts the supply and marginal revenue curves to the right. Over the remainder of the model year, our estimates of leftward shifting demand lowers the shadow value of inventories, resulting in a decline in the optimal price of a vehicle. This leads to an average vintage premium within the model of 7.9% over the year. The market-equilibrium model does slightly over-estimate sales during the first year of the product cycle and under-estimate during the second year of the product cycle. But the model succeeds in replicating the basic pattern of the data.

To isolate the role inventory policy has on the optimal price paths, we resolve the model removing the revenue tax. This counterfactual experiment can be interpreted as allowing the firm to engage in a build-to-order inventory policy rather than a build-to-stock policy. In figure 20 we plot the price and inventory




Figure 20: Simulation of monthly retail prices and inventories under a build-to-order policy.

pathes under a build-to-order policy while still allowing for multiple vintages. To ease comparison, we set the scales of the figure consistent with those in figures 15 and 18. As seen in the figure, the firm sets a downward sloping price path that is far less dramatic than the price path with the revenue tax. Prices fall only \$644 (from \$19,422 to \$18,778) or 3.3 percent over the product cycle. For the first 40 weeks of the product cycle 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 stops reducing prices 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 nine weeks into the following model year with an average vintage premium of 3.3 percent. Hence a model with just declining demand yields within-product-cycle price declines that are only 1/4 of the declines generated by a model with a build-to-stock inventory motive. Interestingly, the firm produces and thus sells more vehicles without the revenue tax (188,000 versus 154,500), as early in the model year (when demand is high) the firm sets a lower price and sells more vehicles.

3 Conclusion

In this paper we have documented a set of stylized facts for the within-model-year pricing and sales of new automobiles. In particular, we find that prices decline steadily over the model year while sales and inventories are hump-shaped. This decline in prices occurs throughout the entire model year; 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!” The process of “slashing prices” happens year-round. 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-quarter these price declines, our model suggests that a driving force behind the magnitude of these declines is inventory management.

Advances in production and information technology have made it easier to implement build-to-order policies. Our analysis suggests that enacting such a policy will dampen within-model-year price declines and reduce the period in which consecutive vintage compete with each other.

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., 1981. Departures from marginal-cost pricing in the American automobile industry: Estimates for 1977-78. *Journal of Econometrics* 17, 201-227.
- [6] Bresnahan, T., Reiss, P., 1985. Dealer and manufacturer margins. *Rand Journal of Economics* 16, 253-268.
- [7] Bresnahan, T., Ramey, V., 1994. Output fluctuations at the plant level. *Quarterly Journal of Economics* 109, 593-624.
- [8] Chan, H., Hall, G., Rust, J., 2005. Price Discrimination in the Steel Market. manuscript, University of Maryland.
- [9] 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.
- [10] Esteban, S., Shum, M., 2004. Durable Goods Oligopoly with Secondary Markets: the Case of Automobiles. manuscript, Johns Hopkins University
- [11] Federgruen, A., Heching A., 1999. Combined pricing and inventory control under uncertainty. *Operations Research* 47, 454-475.
- [12] 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.

- [13] Goldberg, P., 1995. Product differentiation and oligopoly in international markets: The case of the U.S. automobile industry. *Econometrica* 63, 891-951.
- [14] Goldberg, P., 1996. Dealer price discrimination in new car purchases: Evidence from the consumer expenditure survey. *Journal of Political Economy* 104, 622-654.
- [15] 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.
- [16] 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.
- [17] Nevo, A. 2000. A practitioner's guide to estimation of random coefficients logit models of demand. *Journal of Economics & Management Strategy* 9, 513-548.
- [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] Ramey, V., Vine D., 2004. Tracking the source of the decline in GDP volatility: An analysis of the automobile industry. NBER working paper 10384.
- [20] Rust, J. 1985. Stationary equilibrium in a Market for Durable Assets. *Econometrica* 53, 783-806.
- [21] Verboven, F. 1996. International Price Discrimination in the European Car Market. *Rand Journal of Economics* 27, 240-268.
- [22] Whiten, T., 1955. Inventory control and price theory. *Management Science* 2, 61-80.
- [23] Zettlemeyer, F., Scott Morton F., Silva-Risso J., 2001. Cowboys or cowards: why are internet car prices lower? manuscript, Yale School of Management.
- [24] Zettlemeyer, 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.