

Copeland, Adam

Working Paper

The dynamics of automobile expenditures

Staff Report, No. 394

Provided in Cooperation with:
Federal Reserve Bank of New York

Suggested Citation: Copeland, Adam (2009) : The dynamics of automobile expenditures, Staff Report, No. 394, Federal Reserve Bank of New York, New York, NY

This Version is available at:

<https://hdl.handle.net/10419/60903>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Federal Reserve Bank of New York
Staff Reports

The Dynamics of Automobile Expenditures

Adam Copeland

Staff Report no. 394
September 2009

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in the paper are those of the author and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the author.

The Dynamics of Automobile Expenditures

Adam Copeland

Federal Reserve Bank of New York Staff Reports, no. 394

September 2009

JEL classification: D12, C61, L62

Abstract

This paper presents a dynamic model for light motor vehicles. Consumers solve an optimal stopping problem in deciding if they want a new automobile and when in the model year to purchase it. This dynamic approach allows for determining how the mix of consumers evolves over the model year and for measuring consumers' substitution patterns across products and time. I find that temporal substitution is significant, driving consumers' entry into and exit from the market. Through counterfactuals, I show that because consumers will temporarily substitute to a large degree, failure to account for automakers' dynamic pricing strategies results in an inaccurate picture of the return to using pricing incentives. A further finding is that the large price discounts typically offered at the end of the model year result in price discrimination by inducing price-sensitive consumers to delay purchasing new vehicles until the later months of the model year.

Key words: price discrimination, discrete-choice demand estimation, automobiles

Copeland: Federal Reserve Bank of New York (e-mail: adam.copeland@ny.frb.org). The author thanks Ana Aizcorbe, Liran Einav, Ron Goetler, Ryan Greenaway-McGrevy, Kathleen Johnson, and Ahmed Khwaja, as well as seminar participants at Johns Hopkins University and the Federal Reserve Banks of Cleveland and New York, for comments and suggestions. This paper was written while the author was an economist at the Bureau of Economic Analysis. The views expressed in this paper are those of the author and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

A fundamental question in durable goods analysis is the degree to which consumers time their purchase decisions in response to changes in price. Temporal substitution enriches the consumer's problem with dynamics and raises a host of issues outside the usual static framework. With temporal substitution, one-time price discounts may draw consumers away from competing products over several time periods, generating large revenue impacts, or they may simply cannibalize future sales. Further, the existence of consumer dynamics reinforces the importance of accounting for dynamic pricing strategies. In response to a price discount today, competing firms may lower their products' prices in the future, discouraging consumers from switching their purchase decisions and dampening the revenue gains from temporary sales. Finally, the existence of temporal substitution allows for the possibility of price discrimination, if consumers differ in their sensitivity to price.

This paper measures consumers' willingness to time their purchases and examines the issues described above for the new motor vehicle market. Despite being an oft-studied industry, little is known about the degree to which consumers are willing to change the timing of their new automobile purchases. Yet the market for new motor vehicles is a natural place to look for temporal substitution, given that new cars are expensive, firms frequently offer price discounts, and households have easy access to car repair services to prolong the life of their current vehicle. Indeed, the industry wisdom is that consumers frequently "time" their purchase decisions. To explain the jagged-saw profile of light motor vehicle sales (see Figure 2), analysts often rely on temporal substitution, claiming temporary sales have "pulled forward" customers.¹

Suggestive evidence that temporal substitution plays a substantial role in the motor vehicle market is the fall in prices over the model year.² While several factors may cause automakers to adopt this pricing strategy, a leading possibility is that firms are taking advantage of heterogeneous consumers to price discriminate. Supporting this hypothesis, Aizcorbe, Bridgman, and Nalewaik (2007) show that the average income of new vehicle purchasers falls 10 to 15 percent over the model year.

¹Light motor vehicles are those vehicles purchased by households (e.g., cars, pickups, and SUVs). The terms "automobile" and "light motor vehicle" are used interchangeably in this paper.

²Copeland, Dunn, and Hall (2008) document that U.S. light motor vehicle prices fall steadily at an average annual rate of 9 percent.

Past work in this industry has employed static demand models, which are unable to answer questions about temporal substitution.³ This paper builds upon this literature by estimating a dynamic demand model for new vehicles within the model year. To accomplish this, I construct and estimate a dynamic optimal stopping problem, where consumers decide both if they want to purchase a new vehicle and when in the model year to purchase it. Conditional on deciding to purchase, the trade-off a consumer faces is to buy and immediately enjoy a new vehicle, or to wait and purchase that vehicle in the future at a lower price.

Two features of the data allow for measuring consumers' responsiveness to price within a dynamic setting. First, by focusing on the timing of purchases within the model year, I take advantage of a peculiarity of the automobile market whereby firms simultaneously introduce new vintages of their products every fall.⁴ This provides a twelve month window where consumers are facing the same choice set over time, but with varying prices. As such, the timing of consumers' purchases cleanly reflects the trade-off of immediately enjoying a new vehicle versus waiting to purchase the same vehicle later at a lower price. Second, I merge the price and sales data with demographic information, which allows for observing the mean income of new vehicle purchasers over the model year. These data play a crucial role in quantifying to what degree households' sensitivity to price explains the time-series of prices and sales of automobiles.

I find that consumers are sensitive to price and are willing to alter the timing of their purchase decision throughout the model year. This result stands in contrast to the industry wisdom, which believes that consumers are willing to change the timing of their new automobile purchases by at most one or two months. Further, the number of households that temporally substitute in response to a price change dominates the number of households that substitute cross-sectionally or those households that enter or exit the new motor vehicle market. Typically, the number of households temporally substituting from a month close to the period in which the price discount occurs is roughly equal to the number of households that switch within the same month. Hence, if a vehicle is placed on sale halfway through the model year, the number of consumers switching temporally can be five to six times as large as those

³Berry, Levinsohn, and Pakes (1995), Goldberg (1995), and Petrin (2002) are classic papers that employ static models to estimate price elasticities for new motor vehicles in the United States.

⁴Occasionally, firms introduce new vintages at other times of the year, but these are unusual events and such vehicles account for a tiny portion of the market.

switching in the cross section. Finally, the smallest group of switching consumers is represented by those entering or exiting the market in response to a price discount. Accounting for consumers' timing of durable goods purchases, then, is crucial to understanding monthly variations in sales.

Using the model, I examine the impact of a one-time, unexpected \$500 price discount on all General Motors (GM) pickup trucks in the sixth month of the model year. I show that households from all six remaining months of the model year push up the timing of their motor vehicle purchases. The single largest source of switching consumers are those households that would purchase GM vehicles over the remainder of the model year. But these cannibalized sales are more than offset by consumers switching from other brands. Over the entire model year, GM's one-time sale generates an additional 1,452 units in sales, an increase of 0.03 percent. In revenue terms, however, the one-time sale is expensive for GM, resulting in an \$18.7 million loss over the model year, or a decrease of 0.02 percent.

This counterfactual, though, does not account for competing firms' price responses. Because the automobile market is an oligopoly and prices are flexible, we expect firms to react to each other's price discounts. Following Erdem, Imai, and Keane (2003) and Gordon (forthcoming), I use a reduced-form approach to approximate firms' price responses. Incorporating these predictions into the structural model, I estimate revenue losses that are 50 percent larger than the case without price responses. Not accounting for a firm's optimal dynamic pricing behavior, then, substantially understates the revenue losses from placing a vehicle on sale.

Finally, the paper considers what happens to market-level sales and revenues if automakers collude by no longer offering price discounts. I find that, over the entire model year, 0.33 percent fewer units are sold and 3.2 percent more revenue is earned, compared to the case with incentives. Revenues are boosted for two reasons. First, there is a general increase in prices owing to the lack of incentives. Second, because the option value of waiting is now lower due to the lack of incentives, some consumers purchase their vehicles earlier in the model year, when prices are higher. The sizable shift in consumers from later to earlier parts of the model year demonstrates the existence and importance of price discrimination in this market.

A small but growing literature has estimated consumers' temporal price elasticities. Carranza (2003), Carranza (2006), Song and Chintagunta (2003), Gordon (forthcoming), Nair (2007), and Gowrisankaran and Rysman (2007) are recent works that estimate dynamic mod-

els of consumer demand for cameras, video game consoles, computer processors, and DVD players, using a variety of empirical techniques. This paper builds upon this body of work in two ways. First, it considers a new industry, one that is an important sector of the U.S. economy in and of itself. Further, the results have importance beyond motor vehicles, because this industry is often studied by macroeconomists for insights on aggregate inventory behavior (e.g., Hall (2000), Attanasio (2000)), the volatility of GDP (e.g., Ramey and Vine (2006)), exchange rate pass-through (e.g., Goldberg and Verboven (2001)), and other macroeconomic issues. Understanding consumers' willingness to temporally substitute impacts these macroeconomic studies.

Second, relative to past work, this paper's empirical application leans less heavily on assumptions about expectations. In previously studied durable goods markets, both prices and the choice set change from one period to the next. A difficult empirical hurdle, then, is identifying how future price changes, versus the potential introduction of new good, influence a consumer's purchasing decision. Because automakers simultaneously release new model-year vehicles, within the model year I observe consumers' purchases with varying prices, but given a fixed choice set. This allows us to cleanly estimate consumers' temporal price elasticities.

A closely related literature examines the interaction of primary and secondary markets. In particular, Shum and Esteban (2007), Schiraldi (2006) and Chen, Esteban, and Shum (2008) consider dynamic demand models of new and used automobile markets. They focus on the longer-run problem of durable goods replacement and model this industry at the annual frequency. In contrast, this paper looks only at the primary market for motor vehicles and focuses on the shorter horizon where households can temporally substitute within the model year.

This paper also contributes to the literature on understanding various aspects of automobile pricing. In a series of papers, for example, Busse, Silva-Risso, and Zettelmeyer (2006), Zettelmeyer, Morton, and Silva-Risso (2007), and Busse, Simester, and Zettelmeyer (2007) analyze the impact of incentives, dealer inventories, and price cues, respectively, on consumer behavior. This paper adds to this line of research by quantifying consumers' temporal substitution and so provides a understanding of whether different pricing strategies are drawing new customers into the market, pulling consumers away from competing products, or simply changing the timing of existing consumers' purchase decisions.

Finally, this paper complements an existing literature that focuses on frequently purchased

nondurable goods. A number of studies focus on the high/low pricing strategies often used by retail firms. Work by Slade (1998), Aguirregabiria (1999), Pesendorfer (2002), Erdem, Imai, and Keane (2003), and Hendel and Nevo (2006), for example, uses grocery retail data to analyze, among other issues, if temporal substitution drives a firm's pricing strategies. Because these papers focus on frequently purchased nondurable products, the consumer behavior driving these pricing strategies differs from that examined in this paper.

The remainder of the paper describes the data set (Section 2), constructs a structural model (Section 3), and lays out the estimation strategy (Section 4). The results (Section 5), counterfactuals (Section 6), and conclusion (Section 7) follow.

2 Data


The data used covers the vast majority of the light motor vehicle market in the United States. I use a data set on monthly prices and sales by model and model year similar to that in Copeland, Dunn, and Hall (2005). Unlike in that paper, however, I include vehicles produced by foreign manufacturers because the lack of information on overseas production is not a concern here. Further, the panel of data used here is much longer. The sales data come from Wards Communications, while the price data are derived from retail transactions captured at dealerships by J.D. Power and Associates (JDPower).⁵ JDPower attempts to measure precisely the prices customers pay for their vehicle, even adjusting the price when a dealership under- or overvalues a customer's trade-in vehicle as part of a new vehicle sale. The average cash rebate received by consumers on each model and model year was also obtained from JDPower.

The sample covers the period January 1999 to November 2008. Automobile manufacturers typically release new vintages for each model around September of each year; hence the dataset include 1999 through 2008 model years. In addition to the price and sales data, I merged vehicle characteristic data (e.g., horsepower, weight, and length) from *Ward's Automotive Yearbook* (various years), which provides a rich set of observable characteristics for each vehicle.

As mentioned in the introduction, a striking feature of the automobile market is the constant decline in prices throughout the model year. Figure 1 illustrates the decline in prices by model

⁵Corrado, Dunn, and Otoo (forthcoming) use these price data to analyze recent trends in motor vehicle prices.


Figure 1: Laspeyres Price Indexes by Model Year, 2000-2007



year with a Laspeyres price index. Three forces justify falling prices as the optimal price-setting strategy, all of which are incorporate into this demand model. First, firms could be price discriminating and separating consumers by their sensitivity to price. Those consumers not willing to wait for future price declines would buy at the beginning of the model year.

Second, Pashigian, Bowen, and Gould (1995) postulate that automobiles have a fashion component that falls in value the longer a car has been on the market. The fall in vehicle price over the model year may then simply reflect the decline in the fashion component of the vehicle. Alternatively, the price decline could be a function of informational frictions in the secondary market. Used cars are typically priced based on their model year, not their date of production. Hence, automobiles manufactured at the beginning and end of the model year have the same price on the used car market, not taking into account the second-order price adjustments due to mileage. Industry wisdom is that customers take this imperfection in the used car market into account when purchasing a new vehicle and so demand lower prices at the end of the model year. Hence, the price decline over the model year may represent the forgone value of driving services.

Figure 2: Percentage Change in Aggregate Sales



Third, Copeland, Dunn, and Hall (2008) demonstrate that automakers' build-to-stock inventory management style drives firms to adopt their observed pricing strategies. They claim that firm-held inventories, by increasing the variety of choice available to consumers, increase consumer demand.⁶ Hence, automakers set prices relatively high at the beginning of the model year to dampen demand and allow inventories to accumulate. Firms then gradually lower prices over the model year, once inventories have been built up.

Alongside the falling prices, total light vehicle sales are quite volatile. The percent change in aggregate sales follows a jagged pattern, with most positive percent changes followed by negative percent changes (see Figure 2). Using a two-stage least squares approach, I test whether temporal substitution is a candidate for explaining these price and sales time-series. I first de-trend both price and sales data by separately regressing the log of these two variables on a third-order polynomial of the product cycle index and dummy variables for the calendar

⁶Kahn (1987) and Kahn (1992) also find that inventories are productive in generating greater sales at a given price.

Table 1: Statistics on the Absolute Value of the Price Residuals

Freq	Mean	Std Dev	Median	Min	Max
35,296	.02722	0.02867	0.01928	0	0.60360

month and the model year, for each model in the sample.⁷ There is substantial movement of price around its trend. The mean and median of the absolute value of the price residuals are 0.027 and 0.019, respectively, implying that deviations from trend typically average 2 to 3 percent of the price level (see Table 1). This suggests that, even if there is a decline in the value of vehicles over the model year, there is enough price variation around trend to provide opportunities for consumers to time their purchases.

In the second stage, I look for the presence of temporal substitution by running a series of regressions of the change in sales residuals on the change in price residuals. Letting $(\hat{s}_{jt}, \hat{p}_{jt})$ denote the sales and price residuals for vehicle j in period t from the first stage, I estimate,

$$\hat{s}_{jt} - \hat{s}_{j,t-n} = \zeta + \nu (\hat{p}_{jt} - \hat{p}_{j,t-n}), \quad (1)$$

where n is a positive integer. If consumers are willing to intertemporally substitute, then ν should be negative. If prices rise between periods $t - n$ and t (controlling for trends in price), then I would expect sales to fall between these periods (controlling for trends in sales), because consumers will shift their purchases to period $t - n$. I estimate the above regression ten times, for $n = 1, 2, \dots, 10$, to gain a sense of how many months consumers are willing to shift their new vehicle purchases in response to price discounts. Table 2 reports the results, where each row is a separate regression.⁸ The negative coefficient estimates for the price coefficient, ν , support my claim that consumers are timing their purchases of new motor vehicles. Further, these results suggest that consumer price discounts in period t do not just influence sales in period $t + 1$, but rather affect sales over a much longer horizon. Indeed, the biggest effects of a price change on sales are estimated to be four or five months out. This result is counter to the prevailing view in the automobile industry, that temporal substitution is mainly a localized

⁷Excluded from this analysis are the few models in the sample with less than two model years of data.

⁸For the cases where $n > 10$, the estimated coefficients on the price residuals were not statistically significant at the 90 percent level.

Table 2: Coefficient Estimates from Ten Sales Residual Regressions

Periods of comparison	Price coefficient		Number of observations
	estimate	std error	
$t, t - 1$	-0.434	(0.066)	33,566
$t, t - 2$	-0.735	(0.078)	31,851
$t, t - 3$	-1.087	(0.088)	30,135
$t, t - 4$	-1.163	(0.096)	28,419
$t, t - 5$	-1.163	(0.103)	26,704
$t, t - 6$	-1.026	(0.108)	24,992
$t, t - 7$	-0.905	(0.113)	23,282
$t, t - 8$	-0.777	(0.120)	21,577
$t, t - 9$	-0.567	(0.126)	19,878
$t, t - 10$	-0.370	(0.133)	18,185

Note: Let (\hat{s}_t, \hat{p}_t) denote the residuals in period t from a regression of the log of sales and price, respectively, on a cubic model-year trend with calendar year and month dummies. The difference in sales residuals between periods t and $t - n$, $n = \{1, 2, \dots, 10\}$ is the dependent variable in the regression. The independent variables are a constant term and the corresponding difference in price residuals.

force, where consumers shift their purchase decisions by, at most, one or two months.

To better measure the degree of temporal substitution that occurs in the light motor vehicle market, I develop and estimate a dynamic model of consumer behavior.

3 Model

Similar to Berry, Levinsohn, and Pakes (1995), hereafter BLP, I model consumer purchasing behavior at the household level, within a discrete-choice framework. I then construct the market demand system by aggregating over households, who are heterogenous in income and tastes for motor vehicle characteristics. This approach builds upon the BLP framework by explicitly modeling the sequential nature of the consumer's problem. I assume consumers solve a finite-period optimal stopping problem, where they enter each new model year with an income draw and choose if they want to buy a new vehicle and, if so, when in the model year to purchase it.

For each model year, I assume that consumers can purchase only one vehicle, and that the

remainder of a household's income is used to purchase an alternative composite good. The indirect utility derived from choosing an automobile depends on the interaction between a household's characteristics and a product's characteristics. I distinguish between two types of product characteristics: those that are observed by the econometrician (such as size and height), denoted by \mathbf{X} , and those that are unobserved by the econometrician (such as vehicle-service contracts), denoted by ξ . I model the effect of price on consumer i 's indirect utility through a distaste for price term, η_i . Following Berry, Levinsohn, and Pakes (1999) I assume that $\eta_i = \frac{\eta}{y_i}$, where η is a parameter to be estimated and y_i is a draw from the income distribution.

To separate the consumption of the alternative composite good from the timing of the motor vehicle purchase, I assume consumers can use capital markets to perfectly smooth their consumption of the alternative composite good over the model year. I model the perfect-smoothing assumption by discounting the distaste for price term as if it occurs in the last month of the model year. Hence, the consumer i 's indirect utility from purchasing a vehicle $j \in J$ in month $t \in T$ of a model year is given by

$$u_{ijt} = \mathbf{X}_{jt}\gamma + \sum_{s=1}^{11} 1_{C_t=s}\zeta_s + \xi_{jt} - \beta^{T-t} \frac{\eta}{y_i} p_{jt} + \sum_{k=1}^K \sigma_k v_{ik} x_{jk} + \varepsilon_{ijt}, \quad (2)$$

where β is the monthly discount factor, p_j denotes the vehicle price, and $x_{jk} \in \mathbf{X}_j$ is the k th observable characteristic of product j . The indicator variable $1_{x=y}$ is equal to 1 when $x = y$, and C_t tracks the number of months since the start of the model year. Consequently, the dummy variables ζ_s measure the trend in utility over the model year (e.g., declining utility due to fashion). Within this framework, the mean utility of a vehicle, δ_{jt} , is given by $\mathbf{X}_{jt}\gamma + \sum_{s=1}^{11} 1_{C_t=s}\zeta_s + \xi_{jt}$, where γ are parameters to be estimated. Consumers then have a distribution of tastes over the observable characteristics. For each characteristic k , consumer i has a taste v_{ik} , which is drawn from an independently and identically distributed (iid) standard normal distribution. The parameter σ_k captures the variance in consumer tastes. Finally, ε_{ijt} , a taste-for-variety shock, is a random variable that is i.i.d. over consumers, vehicles, and time, and is distributed type 1 extreme value.

This indirect utility function is embedded into a finite-period optimal stopping problem. The consumer maximizes utility by comparing his best choice today, $\operatorname{argmax}_{j \in J} \{u_{ijt}\}$, against

the option value of waiting. In the last month of the model year, this option value is simply the utility value of not purchasing a new vehicle in the current model year. The consumer may choose a range of alternative actions, including waiting to purchase a new car in the next model year, or, alternatively, purchasing a used car. Letting π_{it} denote the utility flow consumer i receives from using his current method of transportation (i.e., the outside option), the continuation value in the last period of the model year can be written as

$$V_{iT} = \max \left\{ \pi_{iT}, \max_{j \in J} \{ u_{ijt} \} \right\}. \quad (3)$$

For the months preceding the last month of the model year, $t = 1, 2, \dots, T - 1$, the consumer's value function is

$$V_{it} = \max \left\{ \pi_{it} + \beta E[V_{i,t+1}], \max_{j \in J} \{ u_{ijt} \} \right\}. \quad (4)$$

The value function lays out the consumer's problem of weighing the utility from purchasing and enjoying a new vehicle now versus waiting. Not surprisingly, the consumer's option value, $V_{i,t+1}$, plays a decisive role in that decision and in the willingness of consumers to temporally substitute.

The expectation over $V_{i,t+1}$ could be taken over the next period's choice set, vehicle characteristics and prices, and ϵ . In this model, however, the expectation is only over ϵ . Given that the choice set and vehicle characteristics do not change over the automotive model year, there is no need for expectations over these variables.⁹ Further, it is assumed that consumers have perfect foresight over prices. Relaxing this assumption imposes large computational costs, especially given the more than 200 vehicles available for purchase every model year. Further, for the purposes of this paper, the gain to incorporating uncertainty over future prices seems small. Owing in part to heavy advertising, consumers seem well aware of the decline in car prices over the model year. Hence, perfect foresight is likely to be a good approximation of consumers' expectations of future prices within the model year.

⁹In this empirical application, one characteristic does change over time: miles per dollar or miles per gallon divided by the price of gasoline.

4 Estimation strategy

To estimate this dynamic model, I mix techniques from the discrete-choice literature with an optimal stopping model, an approach laid out in Gowrisankaran and Rysman (2007).¹⁰ I use a general method of moments (GMM) approach, matching two sets of moments: a set of moments similar to those laid out in BLP and a set of income moments.¹¹

The BLP moments relate to the unobserved characteristic ξ . Except for price, I assume that ξ is uncorrelated with the observed product characteristics of a vehicle, X , or that

$$E[\xi|\mathbf{X}] = 0. \tag{5}$$

For motor vehicles, assuming that product characteristics are predetermined seems reasonable in the short run, because the firm incurs large costs to change these characteristics. The second set of moments is the mean income of households that purchased a new vehicle, by quarter. These data are taken from Aizcorbe, Bridgman, and Nalewaik (2007), who compiled these results from survey data.

To minimize the GMM loss criterion, I use a three-loop estimation routine. The outer-most loop invokes a simplex routine to find the values of (η, σ) that minimize the GMM loss criterion. Given (η, σ) , the estimation algorithm starts with an initial guess of the continuation values of consumers. Based on these guesses, in the inner-most loop of the estimation routine I use the technique detailed in Berry (1994) to back out the mean utility for each vehicle in each month of the model year. Using this new vector of mean utilities, I then move to the middle loop and compute new continuation values through backwards induction. These new value functions are then taken back to the inner-most loop where I compute a new vector of mean utilities. These iterations continue until the difference between iterations for both the vector of mean values and the value functions becomes sufficiently small.¹² Given a final vector of mean utilities that is consistent with a vector of value functions, I can then use ordinary least

¹⁰Gautam Gowrisankaran generously provided the program he used to estimate the dynamic demand model in Gowrisankaran and Rysman (2007), which I used as a guide.

¹¹This approach of adding income moments to the set of BLP-style moments closely follows Petrin (2002).

¹²As noted in Gowrisankaran and Rysman (2007), this system of two loops is not guaranteed to converge. In practice, I experienced nonconvergence problems only for a few parameter vectors that were far from the estimated optimal values.

squares to estimate (γ, ζ) and back out the vector of unobservable characteristics, ξ ,

$$\delta_{jt} = \mathbf{X}_{jt}\gamma + \sum_{s=1}^{11} 1_{C_t=s}\zeta_s + \xi_{jt}. \quad (6)$$

With estimates of ξ , I can compute the set of BLP moments. With the full set of structural parameter estimates, it is straightforward for the model to compute the mean income of households that purchased a vehicle in each quarter. Finally, the GMM loss criterion function can now be calculated, after which the simplex routine, the outer-most loop, chooses a new (η, σ) to evaluate.

An endogeneity problem arises, however, because firms observe ξ and set prices. Hence, ξ is likely to be positively correlated with price, biasing the estimate on the coefficient of price, η . The usual instruments used in the literature are competing products' characteristics. This is because other vehicles' products are correlated with price through equilibrium behavior, but are uncorrelated with ξ because characteristics are assumed to be exogenous.

Unfortunately, product characteristics have little power in this setting because they are constant over the model year. Instead of product characteristics, as instruments I choose monthly dealership inventory levels of competing cars.¹³ The holding of inventories, which is a significant cost to automakers and their dealerships, has been shown to influence the pricing of vehicles (Zettelmeyer, Morton, and Silva-Risso (2007), Copeland, Dunn, and Hall (2008)). Hence, inventories of competing products, through equilibrium pricing behavior, will be correlated with a vehicle's price. Further, the time series of competing products' inventory holdings should be uncorrelated with ξ . In contrast, the inventory holdings and ξ of the *same* vehicle can be expected to be positively correlated. In practice, the instruments used are the sum of inventories of competing vehicles within the same market segment, of competing manufacturers, and of all other vehicles in the market.

Following the literature, I assume all households in the population are in the market for a new vehicle. In the first month of the model year, then, all households consider purchasing a new vehicle. Consistent with the optimal stopping model, in the second month, I determine that only those households that did not purchase a new vehicle in the first month are potential consumers, and so on. I model this shrinking of the number of potential customers by assuming

¹³Model-level inventory data were obtained from Wards Communications.

each consumer i is representative of a type of household that has mass 1. After the first month of the model year, this mass of consumers shrinks by the probability of consumer i buying any new vehicle. This same algorithm is applied after each month in the model year, to properly account for the different purchasing patterns of consumer types. The total number of households is taken from the Current Population Survey's Annual Social and Economic Supplement for years 2000 to 2008. For the years considered, there are a little over 100 million households and a little over 16 million new light-vehicle sales. From this same survey, I also use the distribution of household income. For the results presented here, I used 50 representative consumers.

The identification of the dynamic model's parameters is clearly laid out in Gowrisankaran and Rysman (2007). Intuitively, however, my model is able to measure the consumers' valuations of vehicle characteristics because I observe consumers' purchase decisions over the same choice set while facing different prices.¹⁴ As shown in Section 2, there is substantial variation in price around its trend over the model year. In addition, the set of income moments describing the average income of new-vehicle purchasers plays a large role in pinning down consumers' distaste for price, η . Finally, across model years I observe changes in the choice set alongside corresponding changes in sales.

I do not attempt to estimate the consumer's discount rate, β . As detailed in Rust (1994) and Magnac and Thesmar (2002), the discount rate is rarely identified in these settings. In this environment, the timing of consumer decisions could be a function of impatience or the relative utility value of motor vehicle services versus the alternative composite good. Accordingly, I set consumers' monthly discount rate to 0.9957.

4.1 Data Preparation

To take the model to the data, I had to slightly pare the data. The model is designed to consider consumers' purchase decisions within the model year, not across model years. In the U.S., there is an overlap period in the fall and early winter when the latest two vintages of vehicles are sold simultaneously on the new-car market. I assume that consumers, when shopping for

¹⁴By observing large price changes over time for a given choice set, I avoid the criticism outlined in Akerberg and Rysman (2005), who show that discrete-choice models that rely primarily on changes in the choice set to estimate price elasticities will likely produce biased results.

Table 3: The Effect of Paring the Data

Model year	All Data		Pared Data		Difference	
	Unit Sales (thousands)	Revenue (millions)	Unit Sales (thousands)	Revenue (millions)	Unit Sales (%)	Revenue (%)
2000	17,729	416,414	14,952	352,714	15.7	15.3
2001	16,522	401,868	13,828	335,341	16.3	16.6
2002	17,312	425,032	14,726	363,025	14.9	14.6
2003	16,985	423,494	13,921	350,032	18.0	17.3
2004	17,183	442,485	13,475	348,312	21.6	21.3
2005	16,673	426,214	13,747	354,767	17.5	16.8
2006	16,532	436,643	13,244	347,730	19.9	20.4
2007	15,445	420,451	12,779	346,059	17.3	17.7
Total	134,381	3,392,601	110,672	2,797,980	17.6	17.5

a new vehicle, consider only the newest vintage available. The older vintage is relegated to the outside option. I make this assumption to simplify the consumer's problem and focus on how consumers' evaluate price changes, keeping the choice set constant. Incorporating expectations over changes in the choice set vastly complicates the consumers' problem, with little economic value. As detailed in Table 3, limiting the model year to twelve months eliminated only 17.6 percent of sales and 17.5 percent of revenue. I constructed the twelve-month model year by assuming it started in September of every year. Sales of a particular vintage that occurred before or after this twelve month period were discarded, though the vast majority of the discarded data came from sales after twelve months. Supporting this modeling assumption, Copeland, Dunn, and Hall (2008) find that, within the new light motor vehicle market, cross-price elasticities across vintages of the same model are quite small relative to the cross-price elasticities across the same vintage of different models within the same market segment.

In addition to considering only twelve months of each model year, I drop the 1999 and 2008 model years because I do not observe significant portions of these model years. As such, my data consist of the 2000 through 2007 model years. These eight model years include a total of 386 models and consist of 21,036 observations.

5 Results

This section presents the parameter estimates from the dynamic model. For goodness-of-fit, the results from a comparable static model are also presented. Own-price and cross-price elasticities implied by both models are then presented and compared.

5.1 Parameter estimates

In the empirical specification I choose to include four vehicle characteristics along with a constant term and dummies for each model year in the linear specification of utility, $\mathbf{X}\gamma$ in equation 2. These characteristics are a measure of acceleration (horsepower over weight), miles per dollar (miles per gallon divided by the price of gas per gallon), height, and size (length times width).

For the nonlinear portion of utility, $\sum_{k=1}^K \sigma_k v_{ik} x_{jk}$ in equation 2, I include acceleration, height, size, and a constant term. Adding additional terms, such as miles per dollar, to this nonlinear portion of utility dramatically slowed the estimation algorithm and, in the few instances where a local minimum was eventually found, resulted in own-price and cross-price elasticities quite close to those presented here.

Looking first at the estimated mean values of characteristics, we see that, except for height, the coefficients are all precisely estimated (see the Dynamic Model columns in Table 4). Both the constant and acceleration coefficients are negative and large in absolute value, implying that the mean value to purchasing a new vehicle is negative. These results reflect the data, where only 15 percent of households purchase a new vehicle every model year. The month dummies within the model year are not statistically significant, suggesting that the fall in utility from fashion or the used-car market are offset by the rise in utility associated with increased variety as the model year progresses. The point estimates, though, suggest the increase in variety throughout the model year may be the larger effect.

Turning to the coefficients measuring the heterogeneity of consumers' tastes, I find large differences among consumers. We can infer there is a large variance in the distribution of tastes for acceleration, height, and the utility from purchasing any new vehicle. Reassuringly, heterogeneous tastes for vehicle characteristics are consistent with results from the literature.

Finally, η measures consumers' sensitivity to price, where larger values of η imply an

Table 4: Parameter Estimates

Parameter		Static Model		Dynamic Model	
Mean	Constant	-57.448	(15.050)	-60.286	(1.204)
	Acceleration	0.412	(19.491)	-15.483	(1.937)
	Height	0.716	(52.437)	0.267	(3.736)
	Size	-3.710	(65.127)	7.728	(0.629)
	Miles per Dollar	-1.784	(2.985)	1.388	(0.926)
Month Dummies	1	-0.697	(0.556)	-1.221	(1.122)
	2	-0.228	(0.6031)	-0.659	(1.180)
	3	-0.137	(0.640)	-0.445	(1.195)
	4	0.062	(0.632)	-0.138	(1.181)
	5	-0.109	(0.674)	-0.272	(1.190)
	6	0.074	(0.650)	-0.066	(1.161)
	7	0.315	(0.651)	0.148	(1.171)
	8	0.283	(0.671)	0.080	(1.133)
	9	0.355	(0.690)	0.161	(1.151)
	10	0.293	(0.678)	0.176	(1.138)
	11	0.179	(0.746)	0.136	(1.208)
Variance	Constant	44.453	(11.926)	46.546	(2.322)
	Acceleration	3.420	(2.672)	20.698	(2.863)
	Height	8.418	(3.090)	14.125	(2.975)
	Size	18.031	(1.762)	1.625	(4.811)
Distaste-for-price	η	2.258	(0.428)	1.311	(0.166)
Loss Criterion		21.116		2.515	
Observations		21,073		21,073	

Note: Standard errors are in parenthesis.

increased distaste for price. Given the focus on price elasticities, this parameter is central to our analysis. The parameter is precisely estimated at 1.311 and, as shown in the following section, implies that consumers are somewhat sensitive to changes in price.

To serve both as a benchmark and a measure of robustness, the dynamic model is compared with a standard static model. The static model is similar to the dynamic model except that the market for the static model is the entire model year. To indirectly account for the timing of purchases within the model year, the month of purchase is made a characteristic of the vehicle.¹⁵ Hence, households face an enormous choice set, where, for example, the Ford Focus-January and Ford Focus-February are two different options. Within this static setting, the indirect utility function estimated, letting y denote the model year, is

$$u_{ijy} = \mathbf{X}_{jy}\boldsymbol{\gamma} + \sum_{s=1}^{11} 1_{C_{jy}=s}\zeta_s + \xi_{jy} - \frac{\eta}{y_i} p_{jy} + \sum_{k=1}^K \sigma_k \mathbf{v}_{ik} \mathbf{x}_{jky} + \varepsilon_{ijy}. \quad (7)$$

This equation differs from the dynamic model's indirect utility specification, equation 2, in two respects. First, the number of choices, j , is much larger because there is no longer an explicit time dimension within the model year. Importantly, this implies that consumers know the full utility value of all models throughout the model year. In contrast, in the dynamic model, consumers have expectations over the future utility value of models later in the model year. Second, the price paid for a new vehicle is no longer discounted.

I estimate the static model following the well-known algorithms laid out in BLP and Nevo (2000), using the same instruments and moments described for the dynamic model. Table 4 shows there are a number of similarities between the coefficient estimates of the static and dynamic models. The static model's estimated coefficients also imply that the mean utility from a new vehicle is negative for most households. Further, the static model's estimated coefficients on the month dummies are not statistically significant, but the point estimates suggest an upward trend in utility over the model year, just as in the dynamic model. Finally, the static model finds roughly the same level of heterogeneity in households' tastes for characteristics.

The estimated distaste-for-price coefficient for the static model is twice as large as the dynamic model's coefficient. As shown in the following section, this difference implies that the own-price elasticities in the static model are roughly double those in the dynamic model.

¹⁵I thank Liran Einav for this suggestion.

Table 5: Mean of New Vehicle Purchasers' Income by Quarter

Quarter	Data	Static Model	Dynamic Model
1	74,973	73,565	74,437
2	73,075	72,657	73,570
3	71,460	71,362	71,326
4	68,603	69,974	68,626

Note: All figures are in dollars.

This difference is partly attributable to the lack of any uncertainty in the static model. Unlike the dynamic model, where the taste-for-variety shocks are revealed each period, in the static setting consumers know the utility level from purchasing any vehicle in all months of the model year.

Comparing the measures of fit, we see that the dynamic model's loss criterion is one-tenth the size of the static model's. The dynamic model's superior fit makes a strong case that the static model does a bad job of trying to indirectly capture the dynamic behavior of this market. The static model's worse fit is attributable mainly to the income moments. While the static model fits the general pattern of declining income over the model year, it has a harder time matching the income decline over the model year (see Table 5). The data show a fall of \$6,370 in income from the first quarter to fourth quarter. The dynamic model generates a fall of \$5,811 while the static model performs less well, showing a decline of \$3,591.

5.2 Price elasticities

To measure consumers' price sensitivity, I compute and present own-price and cross-price elasticities, then compare the dynamic model's results with those from the static model as a point of reference.

From the estimated parameters shown in Table 4, I expect, and find, the static model to generate higher own-price elasticities relative to the dynamic model (see Table 6).

It is not straightforward to compare these results with the literature, because past empirical work was done at the annual frequency and relied on list prices or other non-transaction-based price data. Nevertheless, the own-price elasticities from the dynamic model are centered

Table 6: Own-Price Elasticities for the 2002 Model Year (absolute value)

Market Segment	Static Model			Dynamic Model		
	Quartiles			Quartiles		
	1/4	1/2	3/4	1/4	1/2	3/4
CUV	1.31	1.46	1.71	0.85	0.99	1.14
Large car	2.05	2.12	2.17	0.94	0.99	1.05
Luxury car	1.76	1.93	2.12	1.12	1.23	1.42
Midsize car	1.84	1.94	2.03	0.93	1.00	1.11
Pickup	1.75	1.92	2.16	0.74	0.95	1.12
Compact car	1.51	1.64	1.78	0.87	0.96	1.08
SUV	1.38	1.82	2.02	0.72	0.82	1.02
Van	2.26	2.35	2.45	0.53	0.72	0.85
All	1.65	1.89	2.11	0.86	1.02	1.19

around -1, which is less than the results reported in BLP and Goldberg (1995).¹⁶ The results from the static model are closer to the literature, a median own-price elasticity across all market segments of -1.9.

At first glance, the implication that a significant number of products have own-price elasticities that are less than 1 in absolute value is puzzling. Because automakers are profit maximizers and choose price, the estimated own-price elasticities should not be less than 1 in absolute value. In this environment, however, automakers are maximizing profits across both a range of products and over the model year. These own-price elasticities do not account for profit maximization across a product line nor the cannibalization of future sales, given a price change today. Hence, in a dynamic setting where firms sell a range of products, an own-price elasticity that is less than 1 in absolute value does not necessarily imply that the automaker is not maximizing profits. Indeed, in the counterfactuals presented in the next section, the automaker loses substantial revenue over the model year when raising the price of a vehicle in one month.

Turning to the cross-price elasticities, I continue to find that the static model predicts higher amounts of substitution relative to the dynamic model. Table 7 reports aggregate statistics on

¹⁶BLP does not report an average own-price elasticity, but lists a small subset of own-price elasticities that generally range from -6 to -3. Goldberg (1995) reports an average own-price elasticity of -3.28.

Table 7: Within Market Segment Cross-Price Elasticities for the 2002 Model Year

Market Segment	Static Model			Dynamic Model		
	Min	Mean	Max	Min	Mean	Max
CUV	0	0.00037	0.00259	-0.00010	0.00038	0.00801
Large car	0.00008	0.00125	0.00414	0.00002	0.00074	0.00333
Luxury car	0	0.00022	0.00254	-0.00001	0.00022	0.00235
Midsize car	0	0.00123	0.00920	-0.00005	0.00087	0.01015
Pickup	0	0.00325	0.02283	-0.00021	0.00078	0.00747
Compact car	0	0.00227	0.01747	-0.00001	0.00111	0.01156
SUV	0	0.00076	0.00853	-0.00008	0.00036	0.00600
Van	0	0.00172	0.01531	-0.00001	0.00020	0.00300

Note: Statistics are computed only over vehicles within the same market segment.

these elasticities by market segment.

Rather than summarize the elasticities across all vehicles, when constructing the minimum, mean, and maximum cross-price elasticities, I consider only those vehicles within the same market segment. Not surprisingly, including all vehicles in these calculations further lowers the mean cross-price elasticities.

A striking difference between the static and dynamic models is the potential for negative cross-price elasticities within the dynamic model. This possibility reflects the dynamics of the consumer's problem and is partially the reason why the dynamic model fits the data better than the static version. To understand how the dynamic model can generate negative cross-price elasticities, consider the relevant formulas. Letting M denote the size of the market, s denote market share, p denote price, and (j, k) denote vehicles, the cross-price elasticity is defined as

$$\frac{ds_j M}{dp_k} \frac{p_k}{s_j M} = \frac{ds_j}{dp_k} \frac{p_k}{s_j}, \quad (8)$$

where the M s cancel because they are constant over the model year. As is well known, given the expected negative distaste-for-price term, this formula ensures that cross-price elasticities are positive.

In contrast, the dynamic model explicitly models the sequential nature of the consumer's problem. Consequently, the size of the market varies over the year as households purchase a

vehicle and leave the market. Letting t denote a month in the model year, M_t now depends on the choices consumers have seen so far, along with the value of waiting. Taking this into account, the cross-price elasticity in the dynamic model is equal to

$$\frac{d(s_{jt}M_t)}{dp_{k,t+n}} \frac{p_{k,t+n}}{s_{jt}M_t} = \frac{ds_{jt}}{dp_{k,t+n}} \frac{p_{k,t+n}}{s_{jt}} + \frac{dM_t}{dp_{k,t+n}} \frac{p_{k,t+n}}{M_t}, \quad (9)$$

where n is an integer. The first term on the right-hand side, the direct substitution effect, is similar to the price elasticity in the static model (equation 8). The second term, however, is an indirect effect that is not accounted for in the static model. This term captures the effect of price changes on consumers' willingness to wait. Because the taste-for-variety term, ϵ , has unbounded support, consumers have a positive probability of purchasing every vehicle in the choice set. Hence, a change in price of any vehicle today affects overall demand in the current and future periods through the consumer's option value to waiting. A rise in the price of a vehicle in period $t + 1$, for example, lowers the continuation value of period t consumers and so discourages them from waiting. This may result in a lower number of potential consumers available in period $t + 1$ and beyond, dampening demand for all vehicles in the choice set. Consequently, the indirect force allows for the possibility of negative cross-price elasticities. For our parameter estimates, the indirect effect generally causes the dynamic model's cross-price elasticities to be less than the static model's. The indirect effect also generates a small number of negative cross-price elasticities, which drives, at least in part, the dynamic model's better fit of the data relative to the static model.

The computed cross-price elasticities are fairly small, reflecting the large choice set available to consumers. Rather than switch en masse to the nearest competitor, consumers are willing to redistribute themselves across a number of competing vehicles in all months of the model year. Hence, while consumers are most likely to substitute to the closest competitors within the same month of their original purchase, the distribution of switching consumers across all vehicles and months of the model year is diffuse. This result that consumers are not just temporally substituting to and from nearby months, while against industry wisdom, is consistent with the reduced-form results presented in Table 2.

6 Sales and revenue effects of incentives

With the estimated parameters, my dynamic model provides predictions about consumers' dynamic responses to changes in price. Importantly, I can compute how much of a unit sales increase in response to a price decrease comes from temporal substitution versus entry of new consumers into the market. Further, I can decompose the temporal substitution into a cannibalization effect, where consumers are simply switching the timing of their purchase and not their vehicle choice, and a market share effect, where consumers switch from a competitor's product to the vehicle on sale. I present two counterfactuals. The first considers the sales and revenue gains when GM places all of its pickups on sale for one month of the model year. This counterfactual demonstrates the importance of temporal substitution and of strategic pricing behavior in this market. The second counterfactual predicts what would happen to sales and revenues if all automakers colluded by not offering cash-back incentives. This counterfactual shows the importance of price discrimination in the new motor vehicle market.

6.1 The gains from price discounts

I conduct two versions of the first counterfactual, in both cases tracking the changes in unit sales and revenues over the model year. First, I do the usual counterfactual where an automaker introduces an incentive in month t of the model year, holding all other prices constant. Second, I consider the more realistic case that incorporates other automakers' responses to the price discount.

Consider the case where General Motors (GM) places a one-time \$500 price decrease on its pickup trucks in the sixth month of the 2002 model year.¹⁷ Assume that consumers do not anticipate this price discount, but perfectly predict prices over the remainder of the model year. Table 8 shows how this one-time price decrease affects unit sales and revenues over the remainder of the model year.¹⁸ The table lists the sales and revenue changes for GM, Ford, Chrysler, and the group of remaining manufacturers. I highlight the Big Three U.S. automakers

¹⁷Busse, Silva-Risso, and Zettelmeyer (2006) demonstrate that a \$1,000 cash-back incentive offered by a manufacturer does not typically reach the consumer in its entirety. Pass-through varies from 70 to 90 percent because of information asymmetries. We care about the prices consumers pay and so abstract from how much the manufacturer must offer in order to induce a \$500 decline in the actual transaction price.

¹⁸I do not impose any constraints on the supply of vehicles available for sale.

Table 8: \$500 Price Decline for GM Pickups in Month 6

Month	Chrysler	Ford	GM	All others
Unit sales: difference between data and counterfactual				
6	-23.5	-33.2	1,801.0	-55.7
7	-23.9	-35.8	-55.5	-58.6
8	-23.3	-33.7	-55.1	-54.4
9	-25.1	-35.8	-54.3	-58.5
10	-24.0	-35.3	-60.9	-56.0
11	-22.2	-32.6	-63.6	-56.0
12	-23.9	-36.0	-59.8	-61.0
Total	-165.9	-242.3	1,451.7	-400.0
Revenue (thousands of dollars): difference between data and counterfactual				
6	-540.7	-825.7	-9,454.7	-1,449.9
7	-566.2	-885.7	-1,448.9	-1,516.9
8	-534.1	-829.8	-1,414.6	-1,412.2
9	-571.0	-859.0	-1,391.0	-1,492.8
10	-548.4	-832.9	-1,562.1	-1,396.0
11	-521.0	-775.4	-1,726.1	-1,366.0
12	-570.1	-868.6	-1,663.9	-1,459.6
Total	-3,851.5	-5,877.2	-18,661.4	-10,093.3

Table 9: \$500 Price Decline for GM Pickups in Month 6 with Price Responses

Month	Chrysler	Ford	GM	All others
Unit sales: difference between data and counterfactual				
6	-25.6	166.6	1,737.6	-67.6
7	249.4	722.5	-119.9	-109.8
8	256.2	152.5	-116.3	-57.2
9	-56.0	-78.2	-112.7	-133.0
10	-53.5	-77.5	-127.5	-128.2
11	-49.3	-72.8	-130.8	-129.6
12	-53.3	-80.1	-124.2	-142.0
Total	267.9	733.0	1,006.2	-767.5
Revenue (thousands of dollars): difference between data and counterfactual				
6	-1,394.2	-3,828.7	-10,968.7	-3,112.0
7	-3,885.6	-9,041.0	-2,992.2	-3,458.5
8	-3,775.0	-3,685.1	-2,854.2	-3,199.4
9	-1,254.8	-1,836.2	-2,758.1	-3,206.5
10	-1,205.3	-1,794.5	-3,132.0	-3,031.2
11	-1,141.2	-1,694.8	-3,423.2	-3,009.8
12	-1,251.7	-1,896.0	-3,318.4	-3,245.9
Total	-13,907.8	-23,776.3	-29,446.6	-22,263.2

because they dominate the pickup market segment. Because consumers are quite willing to push up their purchase of a new vehicle, the number of consumers temporally substituting is roughly six times as large as those changing their purchase decision within month 6. Indeed, those consumers who originally purchased near the end of the model year are just as willing to temporally substitute as those in month 7. The reason for this result is partly that consumers at the end of the model year have lower incomes relative to those households who purchased earlier in the year, and so they are more price sensitive. Further, only 643 new consumers enter the market in response to the sale, a small fraction of all consumers altering the timing of their purchase decision. The dynamic effects of sales, then, play the primary role in determining the profitability of price discounts.

The price discount by GM does induce a hefty amount of cannibalization. One-fifth of the households that purchased GM pickups in month 6 in response to the price discount would have purchased a GM pickup in months 7 to 12. Overall, however, GM increases its sales by 1,452 units, a 0.03 percent increase in total sales, while revenues fall by \$18.7 million, a 0.02 percent decrease.

Because Chrysler, Ford, and GM dominate the pickup market segment, it is not surprising that Chrysler and Ford lose the most consumers when GM puts its pickups on sale. But the sales and revenue losses by these companies raise the natural question about how automakers respond to one another's price discounts.

With this motivation, we turn to the second case of the first counterfactual, where I account for automakers' price response to GM's price discount. Following Erdem, Imai, and Keane (2003) and Gordon (forthcoming), I use a reduced-form approach to approximate automakers' price-discounting strategies. This paper aims to roughly approximate these strategies. I simplify the strategic pricing game by relying on industry wisdom and assuming, at least for pickups, that GM is a leader in offering incentives.¹⁹ Hence, I model the setting of incentives as a two-stage game. In the first stage, GM sets its price discount for pickups. In the second stage, the remaining manufacturers of pickups set their discounts, in response to GM's incentives. I estimate how much manufacturers respond to GM's incentives with a regression. Denote the sum of the average cash-back received by purchasers of pickup j as d_{jt} , and let

¹⁹Most famously, GM started the 0 percent financing marketing push after the September 11, 2001, terrorist attacks and the "employee discount pricing" program in the summer of 2005. These aggressive marketing moves were quickly followed by Ford and Chrysler and, to a lesser extent, by other automakers.

Table 10: Predicted Incentive Responses to a \$500 GM Incentive in Month 6

	Month			
	6	7	8	9
GM	500	0	0	0
Chrysler	20	237	213	0
Ford	96	259	79	0
All others	50	20	52	0

Note: The \$500 price discount for GM in month 6 is taken as given. The price discounts listed for Chrysler, Ford, and all others are predicted by coefficients estimated by a regression model.

$\bar{d}_{GM,t}$ denote the average incentives received by purchasers of GM pickups. Let c_t equal 1 for Chrysler, 2 for Ford, and 3 for all non-Big Three manufacturers. At the model level, I de-trend beginning-of-period inventories by regressing the stock of inventories held at dealerships on a quadratic model-year time trend and a set of calendar-year dummies. I denote the residuals from this exercise, \hat{inv}_{jt} and include them in the regression to account for aggregate demand shocks. Finally, let v_t index vehicles and m_t denote a linear model-year trend. Then I run the following regression, looking at the pickup market segment and excluding observations on GM pickups,

$$d_{jt} = \rho m_t + \omega \hat{inv}_{jt} + \sum_{i=1}^J I_{v_t=i} \alpha_i + \sum_{k=1}^3 I_{c_t=k} (\nu_k \bar{d}_{GM,t} + \kappa_k \bar{d}_{GM,t-1}) + \zeta_{jt}, \quad (10)$$

where $I_{x=y}$ is an indicator variable equal to 1 when $x = y$ and ζ_{jt} is an i.i.d. error term. This regression measures how each automaker sets its incentives in response to those offered by GM, controlling for vehicle fixed effects and a model-year trend.

The estimated parameters, reported in the appendix, tell an intuitive story: Chrysler and Ford respond aggressively to changes in GM's incentives, while the remaining manufacturers (e.g. Toyota and Nissan) have small responses. I use equation 11 to predict how automakers respond to a \$500 incentive on GM pickups (see Table 10). Chrysler and Ford have different responses: Chrysler offers large cash-back incentives in the two months following GM's price discount, while Ford immediately responds with a \$96 price discount and a larger \$259 price discount the following month. All other manufacturers have a small initial price response and

a small delayed response.²⁰

Incorporating automakers' price responses has several effects (see Table 9). Not surprisingly, the revenue losses from a \$500 incentive for GM are deepened once I include automakers' price responses, increasing by more than half from -\$18.7 to -\$29.4 million. Ford's and Chrysler's incentive responses result in these firms posting gains in total unit sales but losses in revenues. Overall, then, the incorporation of price responses generates even more temporal substitution, and creates an overall shift of consumers toward the U.S. manufacturers. In addition, the number of new consumers entering the market now stands at 1,240, roughly double the amount predicted in the first counterfactual. Notably, however, GM, Ford, and Chrysler all suffer revenue losses despite the increase in unit sales. Indeed, for the model year, the results for the Big Three are quite similar. All three companies slightly increase their unit sales, but suffer revenue losses in the millions of dollars.

These results highlight the importance of strategic behavior in this industry. While estimating a dynamic consumer demand system is essential for understanding consumers' willingness to temporally substitute and for properly measuring price elasticities, combining this model with information about firms' strategic pricing behavior is essential to predicting and understanding the returns to cash-back schemes and other types of sales.

6.2 Price incentives and price discrimination

The second counterfactual considers the case when automakers collude by not offering incentives over the model year. I consider the 2002 model year, where the average cash rebate at the end of the model year is more than 2.5 times the average rebate at the start of the model year (see Table 11). Without cash rebates, prices fall more slowly over the model year. This means that consumers not only face higher prices at every month of the model year, but also that the option value of waiting has decreased as the large cash rebates at the end of the model year are no longer available.

In response to these changes in price, 48,717 fewer vehicles are sold, but \$11.6 billion more revenue is earned. The sales-weighted average price of a vehicle increases 3.5 percent, from

²⁰As shown in the appendix, only the estimated coefficients for Chrysler's one- and two-month delayed response as well as Ford's one-month delayed response to GM's price discount are statistically significant at the 95 percent level.

Table 11: Average Cash Rebate over the 2002 Model Year

Months within the Automotive Model Year											
1	2	3	4	5	6	7	8	9	10	11	12
312	234	251	345	654	805	757	805	861	938	827	846

Table 12: 2002 Model Year Sales and Revenue with and without Incentives

Month	Unit Sales (thousands)				Revenue (\$ million)			
	CF	Data	diff.	cum. diff.	CF	Data	diff.	cum. diff.
1	516	505	10	10	13,270	12,902	368	368
2	1,138	1,117	21	32	29,185	28,379	807	1,175
3	1,115	1,094	22	53	29,137	28,339	798	1,973
4	1,193	1,175	18	71	31,309	30,421	888	2,860
5	1,040	1,038	1	72	26,722	25,868	853	3,714
6	1,228	1,242	-13	59	31,461	30,403	1,058	4,771
7	1,424	1,436	-12	47	36,337	35,145	1,192	5,964
8	1,358	1,374	-16	32	34,420	33,287	1,133	7,097
9	1,410	1,429	-19	13	35,459	34,302	1,158	8,255
10	1,413	1,441	-27	-14	35,392	34,233	1,159	9,413
11	1,392	1,410	-17	-32	35,232	34,140	1,092	10,505
12	1,449	1,466	-17	-49	36,721	35,606	1,116	11,621

note: “CF” stands for counterfactual where automakers collude and do not offer any incentives, “diff.” stands for difference, and “cum. diff.” stands for cumulative difference.

\$24,651 to \$25,525. Given that revenues increase while vehicles sold decrease, automakers’ profits must increase sizably. When automakers do not offer incentives, consumers both leave the market and shift the timing of their motor vehicle purchases to earlier in the model year. Table 12 compares sales and revenue flows in the data against the counterfactual. Without incentives, 71,000 more units are sold within the first four months of the 2002 model year and \$2.8 billion more revenue is generated. Indeed, the model predicts a large movement of consumers from the last six months of the 2002 model year to the first five months.

The sizable shift in consumers from the later to the earlier months of the model year comes about because the lack of incentives has dampened the gains from waiting. Automakers, then, benefit in two ways from agreeing to not offer any incentives. First, the prices of all vehicles

in all months are higher. Second, because the option value to waiting has been reduced, consumers move up the timing of their motor vehicle purchases in the model year and pay even higher prices. Indirectly, this exercise demonstrates that price discrimination plays a powerful role in spreading consumers of varying price sensitivities over the model year.

7 Summary and conclusion

In this paper I estimate a dynamic demand model for light motor vehicles. I show that consumers are price sensitive and are willing to substitute temporally. Further, I demonstrate that temporal substitution dominates consumers' substitution patterns within the model year. In response to a price cut, the number of households that substitute temporally will be several times larger than those switching in the cross-section or those entering the market.

Through a counterfactual, I compute the revenue gains from offering a one-time price discount. Holding all else constant, a \$500 price incentive on GM pickups in the eighth month of the model year generated, over the entire model year, a 0.03 percent increase in unit sales and a 0.02 percent decrease in revenue for GM. These results, however, under-estimate the revenue losses from price cuts because they fail to account for competing firms' price responses. I show that the revenue losses from the \$500 price cut increase by 50 percent once competing firms' strategic pricing responses are incorporated into the counterfactual. Finally, I show that price discrimination is important in the new-vehicle market. If automakers collude to not offer price incentives, which lowers consumers' value to waiting, then sizable numbers of households move up the timing of their new vehicle purchase.

Incorporating dynamics into the demand system for motor vehicles is important in helping to answer a multitude of other questions. In another paper I aim to analyze the cost-of-living price index implied by the model and compare it to existing measures (which are derived from static models of demand). Other research agendas include expanding this dynamic model to consider consumers' substitution patterns across vintages. Even though dealerships sell two vintages of the same model for only a few months of the model year, it should be possible to estimate how willing consumers are to switch from one vintage to another. This, in turn, would provide information about the strength of the connection between primary and secondary markets and to what degree automakers' profitability is impacted by the market for used cars.

8 Appendix

The following table lists the estimated coefficients, except for vehicle fixed effects, from the regression described in equation 11. I denote the sum of the average cash-back received by purchasers of pickup j as d_{jt} , and let $\bar{d}_{GM,t}$ denote the average incentives received by purchasers of GM pickups. Let c_t equal 1 for Chrysler, 2 for Ford, and 3 for all non-Big Three manufacturers. At the model level, beginning-of-period inventories are de-trended by regressing the stock of inventories held at dealerships on a quadratic model-year time trend and a set of calendar-year dummies. I denote the residuals from this exercise as \hat{inv}_{jt} and include them in my regression to account for aggregate demand shocks. Finally, let v_t index vehicles and m_t denote a linear model-year trend. Then I run the following regression, looking at the pickup market segment and excluding observations on GM pickups,

$$d_{jt} = \rho m_t + \omega \hat{inv}_{jt} + \sum_{i=1}^J I_{v_i=j} \alpha_j + \sum_{k=1}^3 I_{c_t=k} (\iota_j \bar{d}_{GM,t} + \kappa_j \bar{d}_{GM,t-1}) + \zeta_{jt}, \quad (11)$$

where $I_{x=y}$ is an indicator variable equal to 1 when $x = y$ and ζ_{jt} is an i.i.d. error term. This regression measures how each automaker sets its incentives in response to those offered by GM, controlling for vehicle fixed effects and a model-year trend. This regression was estimated only for the pickup market segment and excluded observations on GM pickups. There were 768 observations, and the R-squared is 0.485.

Table 13: Estimated Coefficients from Price Discount Regression Model

Parameter description		Estimated	Std err	95% confidence interval
Chrysler	no lag (ι_1)	0.0410	0.1712	(-0.2952, 0.3772)
	1 lag (κ_1)	0.4728	0.2335	(0.0144, 0.9311)
	2 lag (κ_1)	0.4263	0.1626	(0.1070, 0.7458)
Ford	no lag (ι_2)	0.1922	0.1643	(-0.1305, 0.5148)
	1 lag (κ_2)	0.5169	0.2224	(0.0804, 0.9534)
	2 lag (κ_1)	0.1584	0.1536	(-0.1432, 0.4600)
All others	no lag (ι_3)	0.1001	0.1055	(-0.1071, 0.3072)
	1 lag (κ_3)	0.0399	0.1429	(-0.2407, 0.3204)
	2 lag (κ_1)	0.1045	0.1007	(-0.0933, 0.3023)
Model-year trend	ρ	39.9207	8.6568	(22.9261, 56.9153)
Inventories	ω	0.0007	0.0014	(-0.0021, 0.0034)

References

- ACKERBERG, D., AND M. RYSMAN (2005): “Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects,” *RAND Journal of Economics*, 36, 771–788.
- AGUIRREGABIRIA, V. (1999): “The Dynamics of Markups and Inventories in Retailing Firms,” *Review of Economic Studies*, 66, 275–308.
- AIZCORBE, A., B. BRIDGMAN, AND J. NALEWAIK (2007): “The Implications of Heterogeneous Buyers for Measuring Quality Change,” Bureau of Economic Analysis mimeo.
- ATTANASIO, O. (2000): “Consumer Durables and Inertial Behavior: Estimation and Aggregation of (S,s) Rules for Automobile Purchases,” *Review of Economic Studies*, 67(4), 667–696.
- BERRY, S. (1994): “Estimating Discrete Choice Models of Product Differentiation,” *RAND Journal of Economics*, 71, 581–611.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63, 841–890.
- (1999): “Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy,” *American Economic Review*, 89, 400–430.
- BUSSE, M., J. SILVA-RISSO, AND F. ZETTELMEYER (2006): “\$1000 Cash Back: The Pass-Through of Auto Manufacturer Promotions,” *American Economic Review*, 96, 1253–70.
- BUSSE, M., D. SIMESTER, AND F. ZETTELMEYER (2007): “The Best Price You’ll Ever Get: The 2005 Employee Discount Pricing Promotions in the US Automobile Industry,” NBER Working Paper 13140.
- CARRANZA, J. E. (2003): “Product Innovation and Adoption in Market Equilibrium: The Case of Digital Cameras,” University of Wisconsin mimeo.
- (2006): “Estimation of Demand for Differentiated Durable Goods,” University of Wisconsin mimeo.

- CHEN, J., S. ESTEBAN, AND M. SHUM (2008): "How Much Competition Is a Secondary Market?," Johns Hopkins University working paper.
- COPELAND, A., W. DUNN, AND G. HALL (2008): "Prices, Production and Inventories Over the Automobile Model Year," NBER Working Paper 11257.
- CORRADO, C., W. DUNN, AND M. OTOO (forthcoming): *Price and Productivity Measurement*chap. Incentives and Prices for Motor Vehicles: What Has Been Happening in Recent Years? Trafford Press.
- ERDEM, T., S. IMAI, AND M. KEANE (2003): "A Model of Consumer Brand and Quality Choice Dynamics Under Price Uncertainty," *Quantitative Marketing and Economics*, 1, 5–64.
- GOLDBERG, P. K. (1995): "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," *Econometrica*, 63, 891–951.
- GOLDBERG, P. K., AND F. VERBOVEN (2001): "The Evolution of Price Dispersion in the European Car Market," *Review of Economic Studies*, 67, 811–848.
- GORDON, B. (forthcoming): "Estimating a Dynamic Model of Demand with Durable Goods," *Marketing Science*.
- GOWRISANKARAN, G., AND M. RYSMAN (2007): "Dynamics of Consumer Demand for New Durable Goods," University of Arizona mimeo.
- 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.
- HENDEL, I., AND A. NEVO (2006): "Sales and Consumer Inventory," *RAND*, 37(3), 543–561.
- KAHN, J. (1987): "Inventories and the Volatility of Production," *American Economic Review*, 77(4), 667–79.

- (1992): “Why is Production More Volatile than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory-Holding,” *Quarterly Journal of Economics*, 109(3), 565–92.
- MAGNAC, T., AND D. THESMAR (2002): “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, 70, 801–816.
- NAIR, H. S. (2007): “Intertemporal Price Discrimination with Forward-Looking Consumers: Application to the U.S. Market for Console Video-Games,” *Quantitative Marketing and Economics*, 5(3), 239–292.
- NEVO, A. (2000): “A Practitioner’s Guide to Estimation of Random Coefficients Logit Models of Demand,” *Journal of Economics and Management Strategy*, 9, 513–548.
- PASHIGIAN, P., B. BOWEN, AND E. GOULD (1995): “Fashion, Styling, and the Within Season Decline in Automobile Prices,” *Journal of Law and Economics*, 38, 281–310.
- PESENDORFER, M. (2002): “Retail Sales: A Study of Pricing Behavior in Supermarkets,” *Journal of Business*, 75, 33–66.
- PETRIN, A. (2002): “Quantifying the Benefits on New Products: The Case of the Minivan,” *Journal of Political Economy*, 110, 705–729.
- RAMEY, V. A., AND D. J. VINE (2006): “Declining Volatility in the U.S. Automobile Industry,” *American Economic Review*, 96, 1876–1889.
- RUST, J. (1994): *Handbook of Econometrics*, chap. 51, pp. 3081–3143. North Holland, 4 edn.
- SCHIRALDI, P. (2006): “Automobile Replacement: a Dynamic Structural Approach,” London School of Economics mimeo.
- SHUM, M., AND S. ESTEBAN (2007): “Durable Goods Oligopoly with Secondary Markets: the Case of Automobiles,” *RAND*, 38(2), 332–354.
- SLADE, M. E. (1998): “Optimal Pricing with Costly Adjustment: Evidence from Retail-Grocery Prices,” *Review of Economic Studies*, 65, 87–107.

SONG, I., AND P. CHINTAGUNTA (2003): “A Micromodel of New Product Adoption with Heterogenous and Forward-Looking Consumers: An Application to the Digital Camera Category,” *Quantitative Marketing and Economics*, 1, 371–407.

ZETTELMAYER, F., F. S. MORTON, AND J. SILVA-RISSO (2007): “Scarcity Rents in Car Retailing: Evidence from Inventory Fluctuations at Dealerships,” Yale School of Management mimeo.