Intertemporal Substitution and New Car Purchases *

Adam Copeland†

December 12, 2013

Abstract

This article presents a dynamic demand model for motor vehicles. This approach accounts for the change in the mix of consumers over the model year and measures consumers' substitution patterns across products and time. I find intertemporal substitution is significant; consumers are more likely to change the timing of their purchase in reaction to a price increase rather than buy another vehicle in the same period. Further, I find automakers' use of large cash-back rebates at the end of the model year, although boosting overall sales, induces large numbers of consumers to delay their purchases and so pay lower prices.

Key words: discrete-choice demand estimation, automobiles, dynamics
JEL classification: D12, C61, L62

*I thank Victor Aguirregabiria, Ana Aizcorbe, Richard Crump, Liran Einav, Ron Goettler, Gautam Gowrisankaran, Ryan Greenaway-McGrevy, Igal Hendel (editor), Kathleen Johnson, Ahmed Khwaja, Donghoon Lee and Eugenio Miravete as well as two anonymous referees for their comments and suggestions. I started writing this article while at the U.S. Bureau of Economic Analysis. The views expressed here are those of the author and do not necessarily reflect the position of the Bureau of Economic Analysis, the Department of Commerce, the Federal Reserve Bank of New York, or the Federal Reserve System.

†Federal Reserve Bank of New York; e-mail: adam.copeland@ny.frb.org

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. Intertemporal substitution enriches the consumer's problem with dynamics and raises a host of issues outside the usual static framework. With intertemporal 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. Furthermore, a firm's pricing strategies now must take into account how a product's price contour over time impacts the timing of consumers' purchases, and consequently, the flow of revenues.

This article measures the amount of intertemporal substitution in the new light motor vehicle market.[1] 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.[2] Yet the market for new light motor vehicles is a natural place to look for intertemporal 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, industry wisdom is that consumers frequently "time" their purchase decisions. To explain the jagged-saw profile of the growth rate of light motor vehicle sales, analysts often rely on intertemporal substitution, claiming temporary sales have "pulled forward" customers.

I measure consumers' willingness to time their vehicle purchases by using monthly data to estimate a dynamic demand model within the model year. To accomplish this, I construct and estimate an 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 a vehicle later in the model year when prices have fallen.[3]

Two features of the data allow for measuring consumers' responsiveness to price within a

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

[2]Berry, Levinsohn, and Pakes (1995), Goldberg (1995), and Petrin (2002) are classic articles that estimate price elasticities for new motor vehicles in the United States. These articles employ static demand models, and so are unable to answer questions about intertemporal substitution.

[3]Copeland, Dunn, and Hall (2011) document that U.S. light motor vehicle prices fall steadily at an average annual rate of 9 percent.

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 a vehicle from the same choice set later in the model year when prices have fallen. By focusing on the choice within the model year, however, I neglect substitution across model years. Although this substitution choice is captured crudely in the consumer's outside option, not directly accounting for competition across model years could bias the estimates. I present a robustness exercise, however, which partly assuages this concern. Second, I combine the price and sales data with demographic information on the mean income of new vehicle purchasers over the model year. These income data play a central role in estimating households' sensitivity to price.

The parameter estimates reveal substantial heterogeneity in tastes across households, particularly for the taste of buying a new car. The median own-price elasticity is -1.7, and the 25th and 75th percentiles are -1.5 and -1.9, respectively. These estimates are lower than those reported in the literature, but most previous work employed static models and relied upon listed prices. In contrast, I use a dynamic model and have access to price data based on retail transactions. The model reveals an interesting dynamic over the model year whereby the absolute value of the median own-price elasticity increases by 5 percent. This trend reflects the evolution in the mix of households purchasing new vehicles. In particular, high-income (and so price insensitive) households purchase early in the model year because, for them, the option value to waiting and paying a lower price is small. As the model year progresses, the fraction of all purchasing households who are high income declines, driving up the median own-price elasticity (in absolute value).

The implied cross-price elasticities indicate that consumers are quite willing to alter the timing of their new vehicle purchases. Indeed, I find that cross-sectional substitution is small relative to the amount of intertemporal substitution. Consumers, however, do not often radically change the timing of their purchase. In response to an anticipated, temporary price increase, the demand model predicts that on average more than 60 percent of switching house-

[4] Occasionally, 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.

holds alter their timing by 3 months or less.

I then use the estimated demand model to analyze the impact of cash-back rebates in the new vehicle market. Automakers, especially those headquartered in the U.S., have aggressively used these rebates over the sample period. Strikingly, cash-back rebates roughly double in size (both in level and percent terms) between the first and twelfth months of the model year, and so significantly impact the contour of prices. Hence, beyond attracting new consumers into the market or stealing consumers from one another, automakers' use of cash-back rebates impacts the timing of all household's purchase decisions. To isolate and measure this intertemporal effect on households, I use the dynamic demand model to predict sales under a counterfactual with no cash-back rebates, holding all else constant. This comparative static exercise illustrates that cash-back rebates have two major impacts on aggregate sales. By lowering price, they obviously encourage more households to purchase new vehicles over the model year. The rebates, however, also increase every household's option value to waiting and so cause a shift in sales from the beginning to end of the model year. This delay in the timing of purchase costs automakers a substantial amount of revenue.

A small but growing empirical literature has focused on consumers' intertemporal price elasticities. Work by Slade (1998), Pesendorfer (2002), and Hendel and Nevo (2006a) for example, look at retail grocery prices and present evidence that consumers time their purchase decisions. Hendel and Nevo (2006b) and Gowrisankaran and Rysman (2012) estimate structural dynamic demand models and, among other results, demonstrate that static demand models may generate biased estimates of own-price elasticities.[5] Most recently, Hendel and Nevo (2013) estimate consumers' intertemporal price elasticities and use them to empirically study the role of intertemporal price discrimination in storable goods markets. This article builds upon this small body of work. 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 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 intertemporally substitute impacts these macroeconomic studies.

[5]There is also a growing literature in Marketing looking at similar issues, see for example Erdem, Imai, and Keane (2003).

A closely related literature examines the interaction of primary and secondary markets. In particular, Shum and Esteban (2007), Schiraldi (2011) and Chen, Esteban, and Shum (2013) 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 article looks only at the primary market for motor vehicles and focuses on the shorter horizon where households intertemporally substitute within the model year.

This article also contributes to the literature on understanding various aspects of automobile pricing. In a series of articles, for example, Busse, Silva-Risso, and Zettelmeyer (2006), Zettelmeyer, Morton, and Silva-Risso (2007), and Busse, Simester, and Zettelmeyer (2010) analyze the impact of incentives, dealer inventories, and price cues, respectively, on consumer behavior. This article adds to this line of research by quantifying consumers' intertemporal substitution and so provides an 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. In addition, in related articles, Copeland and Hall (2011) and Copeland, Dunn, and Hall (2011) consider how temporary changes in demand and automakers' build-to-stock inventory management policy, respectively, impact pricing. Although automakers solved a dynamic problem, both articles relied on static demand systems. This article complements those two articles by exploring how accounting for dynamics on the consumer side impacts the estimated price elasticities. Unlike in those two articles, here I do not explicitly account for the effect of automakers' inventory on demand. Instead, I use monthly dummy variables to account for the average effect of inventories on demand over the model year.

2 Data

Before describing the dynamic demand model, I introduce the 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 (2011). Unlike in that article, 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 covers a longer time period. The sales data come from Wards Communications and the price data are derived from retail transactions captured at dealerships by J.D. Power and Associates (JDPA).[6] JDPA 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 JDPA.[7] To account for these financing incentives, I set the retail price of a vehicle equal to its average price minus the average cash rebate received by consumers.

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 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 year with a Laspeyres price index. Three forces justify falling prices as the optimal price-setting strategy, all of which are incorporated 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. Supporting this hypothesis, Aizcorbe, Bridgman, and Nalewaik (2010) show that the average income of new vehicle purchasers falls 10 to 15 percent over the model year.

Second, 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.[8]

[6] Corrado, Dunn, and Otoo (2006) use these price data to analyze trends in motor vehicle prices.

[7] For additional details about the data, see Copeland, Dunn, and Hall (2011).

[8] Pashigian, Bowen, and Gould (1995) present an alternative hypothesis that fashion drives, at least in part, the decline in price over the model year.

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

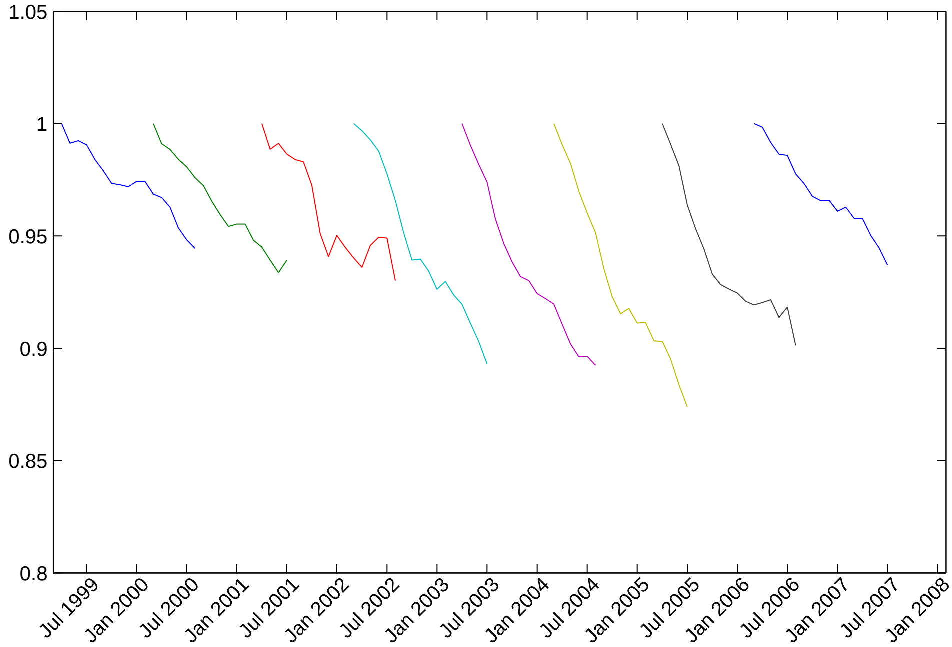

Third, Copeland, Dunn, and Hall (2011) find that automakers' build-to-stock inventory management policy explains four-tenths of the observed pricing strategies. They claim that firm-held inventories, by increasing the variety of choice available to consumers, increase consumer demand.[9] Hence, automakers 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.

The first story about falling prices, price discrimination, relies on the ability of consumers to time their vehicle purchases, whereas the informational frictions and inventory management stories could be consistent with static demand. Falling prices alone, then, do not provide convincing evidence of dynamic demand. To better measure the degree of intertemporal substitution that occurs in the light motor vehicle market, I develop and estimate a dynamic model of consumer behavior.

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

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 in tastes for vehicle characteristics. This approach builds upon the BLP framework by explicitly modeling the sequential nature of the consumer's problem, following the method laid out in Gowrisankaran and Rysman (2012). This model differs from Gowrisankaran and Rysman (2012) in three substantial ways. First, consumers solve a finite-period optimal stopping problem. Second, as made clear later in this section, I assume consumers have perfect foresight over prices, whereas Gowrisankaran and Rsyman assume consumers' expectations over the value of purchasing next period can be approximated by a one-dimensional Markov process. Third, whereas I assume that consumers leave the market and do not return, Gowrisankaran and Rsyman allow for repeat purchases by consumers over time.

Consumers solve a finite-period problem, entering each new model year and deciding if they want to buy a new vehicle and, if so, when in the model year to purchase it. Consumers do not consider substitution across model years, hence the finite nature of the problem. 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 a non-motor vehicle composite consumption good.

A consumer i purchasing a vehicle j at time t receives a net flow of indirect utility at time t equal to

$$f_{ijt} - \alpha_i p_{jt} + \varepsilon_{ijt}, \tag{1}$$

where f_{ijt} is the per period utility flow from using a vehicle j purchased at time t. Vehicle price is given by p_{jt} and α_i denotes consumer i's distaste for price. This distaste-for-price term reflects the tradeoff consumers face in that higher vehicle prices decrease the amount of the composite consumption good a consumer can purchase. Finally, ε_{ijt}, a taste-for-variety shock, is a random variable which captures variations in the purchase experience which do not persist over time.

In equation 1, only f persists over time. Letting $\delta_{ijt} = \sum_{s=t}^{\infty} \beta^{s-t} f_{ijt}$, where β is the monthly discount factor, we can write the consumer's indirect utility from purchasing a vehicle j at time

t in the familiar random coefficients discrete choice framework of BLP,

$$u_{ijt} = \delta_{ijt} - \alpha_i p_{jt} + \varepsilon_{ijt}. \tag{2}$$

I then make δ_{ijt} depend 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 horsepower), denoted by $\mathbf{X}$, and those that are unobserved by the econometrician (such as vehicle-service contracts), denoted by ξ. Following Berry, Levinsohn, and Pakes (1999) I assume that $\alpha_i = \frac{\alpha}{y_i}$, where α is a parameter to be estimated and y_i is a draw from the income distribution. Substituting in for δ_{ijt} I assume the consumer's indirect utility is given by

$$u_{ijt} = \mathbf{X}_{jt}\gamma + \sum_{s=1}^{11} 1_{C_t=s}\zeta_s + \xi_{jt} - \frac{\alpha}{y_i} p_{jt} + \sum_{k=1}^{K} \sigma_k \nu_{ik} x_{jkt} + \varepsilon_{ijt}, \tag{3}$$

where $x_{jkt} \in \mathbf{X}_{jt}$ is the kth observable characteristic of product j at time t. Consumers have a distribution of tastes over the observable characteristics. For each characteristic k, consumer i has a taste ν_{ik}, which is drawn from an independently and identically distribute (iid) standard normal distribution. The parameter σ_k captures the variance in consumers' tastes. 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. Finally, we make the usual assumption that ε_{ijt} is iid over consumers, vehicles, and time, and is distributed type 1 extreme value.

The coefficients on the monthly dummies (ζ_s) capture variation in the consumer's indirect utility from purchasing a new vehicle over the model year, holding everything else fixed. There are two economic forces that drive this variation in indirect utility over the model year.

First, as discussed in section 2, the amount of firm-held inventories (i.e. vehicles on dealership lots) varies throughout the model year. High levels of inventories imply that dealerships have substantial variety of vehicles on their lots, and so they can deliver the vehicle a consumer wants quickly. With low inventories, consumers may experience longer delivery times because dealerships will need to search for the exact model preferred by consumers. This search usually entails locating and acquiring the exact model at another dealership. Hence, variation in firm-held inventories impacts consumers' indirect utility through delivery times. Because automakers tend to introduce new models at the same time, the pattern of firm-held inventories looks the same across most models (Copeland, Dunn, and Hall 2011). Consequently, constraining the coefficients on the monthly dummies to be the same across all vehicles is not unreasonable.

Second, a consumer's indirect utility from a vehicle may steadily decrease over the model year due to an informational friction in the used vehicle market. This informational friction is that prices in the used car market are determined by make and model year and, specifically, not by when the vehicle was purchased and driven off the lot (at which time it starts to physically depreciate). A car sold in January, then, has the same re-sale value as the same car sold in March, despite the January car being two months older. A consumer's indirect utility for a durable good should account for the resale value. Monthly dummies over the model year pick up this decline in the consumer's indirect utility, reflecting that with respect to the re-sale value, a consumer who purchases later in the model year is worse off relative to a consumer who purchases in the beginning of the model year, holding everything else constant.

The specification of indirect utility does not allow me to estimate the impact of these two forces separately. Rather, the monthly dummies capture the net effect of these two forces over the model year.[10]

The indirect utility function is embedded into a finite-period optimal stopping problem. The consumer maximizes utility by comparing his best choice today, $\text{argmax}_{j\in J}\left\{u_{ijt}\right\}$ where J is the set of available vehicles to purchase, 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. Let ε_{i0t} denote the utility flow consumer i receives from the outside option, where ε_{i0t} is an iid random variable which is distributed type 1 extreme value. The continuation value in the last period of the model year is

$$V_{iT} = \max\left\{\varepsilon_{i0T}, \max_{j\in J}\left\{u_{ijT}\right\}\right\}. \tag{4}$$

For the months preceding the last month of the model year, $t = 1,2,\ldots T-1$, the con-

[10]Copeland, Dunn, and Hall (2011) is able to separately estimate the impact of these forces, by using data on firm-held inventories.

sumer's value function is

$$V_{it} = \max\left\{\varepsilon_{i0t} + \beta E[V_{i,t+1}], \max_{j\in J}\left\{u_{ijt}\right\}\right\}. \tag{5}$$

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 critical role in that decision and in the willingness of consumers to intertemporally substitute.

The expectation over $V_{i,t+1}$ could be taken over the next period's choice set, vehicle characteristics, the monthly dummies, 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. An exception is miles per dollar (miles per gallon divided by the price of gasoline), which varies over the model year as gasoline prices fluctuate. I assume that consumers have perfect foresight over fuel efficiency. Further, I assume that consumers have perfect foresight over the monthly dummies and prices.

As opposed to relying on perfect foresight, an alternative approach is to allow for uncertainty in consumers' expectations. Gowrisankaran and Rysman (2012) present a tractable way to implement such an approach, although their method requires the assumption that consumers' expectations over the value of purchasing next period are well-approximated by a one-dimensional Markov process. For the purposes of this article, the gain to incorporating uncertainty over the monthly dummies and future prices do not seem to outweigh the costs. Owing in part to heavy advertising, consumers seem well aware of the decline in car prices over the model year. Consumers are also likely to know that dealerships have relatively small inventories of a model at the beginning of the model year, implying longer delivery times. Hence, perfect foresight is likely to be a good approximation of consumers' expectations of monthly dummies and future prices over the model year.

4 Estimation

In this section I first describe the method used to estimate the model. Then, I detail how the data are parsed to fit the scope of the model.

Method of estimation

The moments

My estimation approach resembles the generalized method of moments (GMM) approach taken by Petrin (2002), in that I supplement the usual BLP-style moments with micro moments.

The BLP moments relate to the unobserved characteristic ξ (see equation 3). I assume that ξ is uncorrelated with the observed product characteristics of a vehicle, $\mathbf{X}$, or that

$$E[\xi(\theta)|\mathbf{X}] = 0, \tag{6}$$

where θ is the vector of parameters to be estimated. For motor vehicles, assuming that product characteristics (except for price) are predetermined seems reasonable in the short run, because automakers incur large costs to change these characteristics.

The usual endogeneity problem arises because firms observe ξ and set prices. Hence, ξ is likely to be positively correlated with price, biasing the distaste-for-price coefficient, α, towards zero. The usual instruments used in the literature are functions of 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.

As instruments, other vehicles' product characteristics have power across model years because of annual changes in vehicle characteristics and in the set of vehicles offered. Product characteristics have no power within the model year, however, because characteristics are constant over the model year. An exception is miles per dollar, which varies over the model year because of changes in the real retail price of gasoline. As an instrument then, competing products' miles per dollar has the advantages of being uncorrelated with ξ and providing price variation over the model year. The disadvantage of this instrument is that relative instrumented prices will not vary over the model year because miles per dollar varies proportionally to the price of gasoline.

Suggesting that miles per dollar is a good instrument, in the data there is a correlation of -0.41 between miles per dollar and vehicle price. Furthermore, recent results in the literature find that automakers adjust vehicle prices in reaction to gasoline prices.[11] In practice, the

[11] Langer and Miller (forthcoming) demonstrate that automakers set prices as if consumers respond to gaso-

instruments I use for vehicle j are the sum of characteristics for vehicles other than j that are in (i) the same market segment, (ii) the same company, (iii) and overall, as well as the vector of vehicle j's characteristics and the monthly dummies.[12] To demonstrate the impact of these instruments, I estimate a static logit version of this model with and without instruments. The non-instrumented estimate of distaste-for-price is -45.7, and the instrumented estimate is -77.1.

The second set of moments, the micro-moments, are the mean income of households that purchased a new vehicle, by quarter. These data are taken from Aizcorbe, Bridgman, and Nalewaik (2010), who compiled these results from survey data. Let $\hat{I}(\theta)$ be a four element vector denoting the model's prediction of the average income of new vehicle purchasers for every quarter of the model year, and let $\mathbf{I}$ denote the average income of new-vehicle purchasers by quarter observed in the data. The second set of moments are

$$E[\hat{I}(\theta)] - \mathbf{I} = 0. \tag{7}$$

Letting $(G_1(\theta), G_2(\theta))$ denote the BLP and income moments respectively, the GMM loss criterion function is

$$\begin{bmatrix} G_1(\theta) \\ G_2(\theta) \end{bmatrix}' \mathbf{W}^{-1} \begin{bmatrix} G_1(\theta) \\ G_2(\theta) \end{bmatrix} \tag{8}$$

where $\mathbf{W}$ is a block diagonal matrix. The block of elements of $\mathbf{W}$ associated with G_1 are set equal to $\mathbf{Z}'\xi(\hat{\theta})\xi(\hat{\theta})'\mathbf{Z}$, where $\mathbf{Z}$ are the instruments and $\hat{\theta}$ is a consistent estimate of the true θ. The block of elements of $\mathbf{W}$ associated with G_2 are set equal to a diagonal matrix, where the elements of the diagonal are equal to the standard error of the mean income of new vehicle purchasers by quarter. The standard errors (taken from Aizcorbe, Bridgman, and Nalewaik (2010)) are supplemental information used to provide relative weights across the income moments. (See the on-line appendix for details on the construction of $\mathbf{W}$.)

The estimation algorithm

line prices, finding that vehicle prices generally decline in the gasoline price. Similarly, Busse, Knittel, and Zettelmeyer (2013) find that changes in gasoline prices are associated with changes in the relative prices of new cars of differing fuel economies.

[12]The vehicle characteristics used to construct the instruments are miles per dollar, horsepower, height and size.

To minimize the GMM loss criterion, I use a simple variation of the approach laid out in Gowrisankaran and Rysman (2012).[13] Because the on-line appendix provides the details of this routine, I provide only an overview in this article.

Given the parameters (α, σ), and random draws of household tastes (ν) and income (y), the GMM loss criterion can be computed. The outer loop of the estimation procedure is a routine that searches over combinations of (α, σ) with the goal of minimizing the GMM loss criterion. The main difficulty to evaluating the loss criterion is recovering the unobserved vehicle characteristic, ξ.

Given consumers' value to waiting, ($\mathbf{V}$, as defined in equations 4 and 5), I can use the contraction mapping described in BLP to recover the values of ξ which equate the level of sales predicted by the model to those observed in the data. Given ξ, I can compute consumers' value to waiting. Unfortunately, I cannot simultaneously solve for ξ and $\mathbf{V}$.

To compute both of these objects so that they are consistent with one another, I employ the following inner loop. I guess initial values for $\mathbf{V}$ and then use the BLP contraction mapping to recover ξ. I then use ξ and equations 4 and 5 to compute, via backwards induction, updated values for $\mathbf{V}$. With the updated $\mathbf{V}$, I again use the BLP contraction mapping to recover new estimates of ξ, and so on.

This inner loop of sequentially updating ξ and $\mathbf{V}$ is continued until there is convergence. Convergence occurs when the average difference between successive iterations of $\mathbf{V}$ is less than 1e-12. When this condition is met, the most recent values of ξ and $\mathbf{V}$ are used to compute the GMM loss criterion.

As noted in Gowrisankaran and Rysman (2012), the inner loop described above is not guaranteed to converge. In practice, I experienced non-convergence problems for only a few parameter vectors that were far from the estimated optimal values.

Following the literature, I assume all households in the population are in the market for a new vehicle. 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. From this same survey, I also use the distribution of household income.

Identification

[13] Gautam Gowrisankaran generously provided the program he used to estimate the dynamic demand model in Gowrisankaran and Rysman (2012), which I used as a guide.

The identification of the dynamic model's parameters is described in Gowrisankaran and Rysman (2012). Intuitively, there are two sources of variation that drive identification. First, across model years I observe both changes in the choice set and in vehicle characteristics alongside corresponding changes in sales. This is the usual source of identification for these types of discrete-choice models. Second, within the model year, I observe consumers' purchase decisions over the same choice set while facing different prices. A caveat with this source of identification is that the instruments for price, while varying over the model year, do not provide variation in relative prices.

Consequently, an important additional source of identification is the average income of new-vehicle purchasers over the model year. These data, captured in the second set of moments described earlier, play a large role in pinning down consumers' distaste-for-price, α. This is because the choice of α drives the model's predictions of the average income of purchasers over the model year. Higher values of α, holding all else constant, generates both a higher average income of purchasers as well as a steeper fall in purchasers' income over the model year. This comes about because increasing α makes consumers become more price sensitive. Relative to those with high incomes, lower income consumers are both more likely to switch to the outside good (raising the average income of new vehicle buyers) and more willing to delay the purchase of their vehicle (creating a steeper fall in purchasers' income over the model year). Conversely, lower values of α generates lower average income of purchasers and less of a fall in income over the model year.[14]

I do not attempt to estimate consumers' discount rate, β. As detailed in 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 $0.95^{1/12}$.

[14] By observing price changes over time for a given choice set and including the set of income moments, I avoid the criticism outlined in Ackerberg 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 1: 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

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 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 consumer's problem. As detailed in Table 1, 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 article's modeling assumption of not considering substitution across model years, Copeland, Dunn, and Hall (2011) 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.

Further, I use ordinary least squares to estimate if, for a given vehicle, changes in old model year prices effect changes in new model year sales, controlling for changes in new model year prices (see appendix A for details). Overall, this approach yields little support for the idea that old model year prices influence new model year sales. Finally, I test the robustness of the structural model with respect to this assumption and find support for this modeling assumption (see section 5). Nevertheless, putting aside the computational restrictions, it would be ideal to include all the data and estimate the substitution patterns both within a model year and across model years.

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, followed by an analysis of the own-price and cross-price elasticities implied by the model. Finally, the robustness of the results are discussed.

Parameter estimates

In the empirical specification for the linear portion of utility, I choose to include four vehicle characteristics along with a constant term, dummy variables for months in the model year, and dummy variables for each model year. The four characteristics are a measure of acceleration (horsepower over weight), height, size (length times width), and miles per dollar (i.e., miles per gallon divided by the price of gas per gallon), which is a measure of fuel efficiency.

For the nonlinear portion of utility, I include a constant term, acceleration, size, and a half-year dummy variable. The dummy variable is equal to 1 for months 7 through 12 of the model year, and equal to 0 otherwise. This variable allows me to estimate the variance on a half-year utility shock (where this shock has a mean of zero). Although ad hoc, the value of this half-year dummy variable is that it allows consumers' tastes for purchasing to vary between the first and second halves of the model year along a dimension not directly tied to price. Without this random coefficient, households' price sensitivity (α/y_i) is the only mechanism which differentially impacts the timing of household's new vehicle purchase. In reality, non-price factors are likely to impact when households purchase their new vehicles. This random coefficient allows the empirical model to account for such effects, albeit roughly. Consequently, the inclusion of the half-year dummy is beneficial because we are then more confident that the estimate of distate-for-price is driven by household's reactions to price differences in the cross-section and over time, rather than to other forces unrelated to prices. A disadvantage of the half-year coefficient is its reduced-form nature, which makes it hard to interpret.[15] In the robustness section (section 5), I further examine the impact of this random coefficient on the estimates.

An improvement over a random coefficient on a half-year dummy would be random coefficients on monthly dummies. With this richer specification we would be even more confident in the estimated coefficient for distaste-for-price and so in the model's predictions about how consumers react to temporary changes in price. I did not use this richer specification because I could not find stable parameters estimates with such a large number of random coefficients. Instead of monthly dummies, an alternative approach would be to assume a distribution for the monthly random preference shock and then estimate the parameters of that distribution. Although this approach only requires estimating few more parameters relative to the specification presented in this article, I did not use this approach because of its computational difficulty.

I limited the nonlinear portion of utility to only 4 terms. Adding additional terms, such as miles per dollar, 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 all these coefficients are precisely estimated (see table 2). The constant coefficient is negative and large in absolute value, driving down the utility value from purchasing a new vehicle. This result reflects the data, where on average only 15 percent of households purchase a new vehicle in a model year. Inline with past work, I find that households value more acceleration, height, and size. In contrast to most past work, I find households value more fuel efficiency.[16] How-

[15]The identification of this coefficient differs from the other random coefficients. In particular, the half-year coefficient is identified by timing within the model year.

[16]In past work, the finding that households do not value fuel efficiency is attributed to fuel efficiency being neg-

Table 2: Parameter Estimates

Parameter		Estimate	Standard Error
Mean	Constant	-23.53	(0.37)
	Acceleration	1.71	(0.38)
	Height	3.61	(0.73)
	Size	4.35	(1.05)
	Miles per dollar	3.73	(0.44)
Month Dummies	1	-0.12	(0.06)
	2	0.32	(0.05)
	3	0.43	(0.06)
	4	0.62	(0.06)
	5	0.48	(0.06)
	6	0.70	(0.05)
	7	0.14	(0.03)
	8	0.19	(0.02)
	9	0.31	(0.02)
	10	0.26	(0.02)
	11	0.16	(0.02)
Variance	Constant	12.83	(0.12)
	Acceleration	3.93	(0.47)
	Size	1.07	(0.22)
	Half-year dummy	4.35	(0.09)
Distaste-for-price	α	2.72	(0.24)
Loss Criterion		496.00	
Observations		21,036	

Note: Coefficient estimates of the model-year dummies are not reported.

ever, although I estimate the expected positive valuation on fuel efficiency, this result is not especially robust; in earlier specifications of this model the estimate on fuel efficiency was negative. What are robust, and are the focus of this article, are the estimated elasticities.

The estimated parameters associated with the monthly dummies indicate there is a rise in the indirect utility from purchasing a car over the first 6 months of the model year. This trend does not carry over into the second half of the model year where the estimated parameters are lower and roughly flat. Recall that these coefficients capture the average net effect of two forces. First, we expect indirect utility to be increasing in the variety of vehicles offered. Second, we expect there to be a decline in indirect utility over the model year because of informational frictions in the secondary market. The pattern of estimated coefficients is consistent with indirect utility increasing over the first 6 months of the model year as automakers rapidly ramp up their inventories (and hence increase the variety available to consumers). In the second half of the year, the lack of a trend may indicate that the two forces effecting the trend in utility over the model year roughly offset one another.

Turning to the coefficients measuring the heterogeneity of consumers' tastes, I find large differences among consumers, results that are consistent with the literature. There is a huge variance on the constant term, implying that people vary widely in their tastes for purchasing a new car. There are also large variances in the distribution of the tastes to purchasing a car in the second half of the model-year and in consumers' tastes for acceleration.

Finally, α measures consumers' sensitivity to price, where larger values of α imply an increased distaste-for-price. Given the focus on price elasticities, this parameter is central to our analysis. The parameter is precisely estimated at 2.72 and, as discussed further in the following section, implies that the median own-price elasticity is -1.7.

Given the estimated parameters, the model is able to explain three-quarters of the income decline by new vehicle purchasers over the model year (see table 3). The data show a fall of \$6,370, or 8.5 percent, in income from the first quarter to fourth quarter of the model year, whereas the model generates a fall of \$4,796, or 6.4 percent. The model, then, does a reasonable job matching the sorting of households by income over the model year. There are two mechanisms in the model contributing to this result. First, households trade-off enjoying a new vehicle early in the model year and paying a high price versus purchasing a new vehicle later

atively correlated with unobserved characteristics of vehicles, especially those related to a vehicle's performance (see Gramlich (2010) for more details).

Table 3: New Vehicle Purchasers' Mean Income by Quarter (dollars)

Quarter	Data	Model
1	74,973	74,454
2	73,075	74,139
3	71,460	70,231
4	68,603	69,658

Note: Quarters are defined over the automotive model year, where the first quarter is composed of September, October and November. The remaining quarters are defined accordingly.

in the model year when prices have fallen. Higher-income households' utility, by construction, are impacted less by higher prices, and so they are more likely, everything else equal, to buy earlier in the model year when prices are high. Second, households differ in their tastes for purchasing in the first and second halves of the model year. For high income households, having a high or low taste regarding the purchase of a vehicle in the second half of the model year most likely impacts the timing of their purchase, not if the household will purchase a new vehicle. As such, these tastes shock will shift when high income household purchase both earlier and later in the model year. In contrast, for low income households these taste shocks are more likely to effect if the household purchases a new vehicle at all. The combination of low prices and high indirect utility from purchasing in the second half of the model year entices some low income households to substitute from the outside option to a new vehicle. Importantly, having a negative and large in absolute value taste for purchasing in the second half is less likely to induce a low income household to purchase a new vehicle, because prices are high at the beginning of the model year. This asymmetric effect of the half-year random coefficient on low income household's purchase decisions helps the model match the income moment.

Price elasticities

To measure consumers' price sensitivity, I compute and present own and cross-price elasticities. These elasticities are based on foreseen, temporary price changes, hence an elasticity of -2 implies that a 1 percent temporary increase in price for model j in month t caused units sold of model j in month t to fall 2 percent. By foreseen, I mean that all consumers at the start of the model year have perfect foresight over future prices.

These elasticities were computed by simulation. To compute the own-price elasticity for model j in month t, I take the observed price path for model j over all months and create an alternative price path. This alternative price path is equal to the observed series except for the month t price which is 1 percent greater than the observed price. I then simulate the model with the alternative price path, and use the difference between simulated and observed sales to compute the own-price elasticity of model j at time t.

The median own-price elasticity implied by the model is -1.69 and the 25th and 75h percentiles are -1.51 and -1.92 (see the last row of table 4). There is variation in elasticities across types of cars, with compact cars, as expected, having the highest median own-price elasticities in absolute value (see table 4).

Highlighting the change in elasticities over the model year, table 5 reports the 25th, 50th, and 75th percentiles of the distribution of own-price elasticities by month in the model year. These elasticities are increasing in absolute value over the model year, with the median elasticity increasing 5 percent (from 1.65 to 1.73) from the first to last month of the model year. The distribution of elasticities also gets narrower over the model year, with the 25th percentile increasing from 1.43 to 1.58 whereas the 75th percentile remains roughly constant around 1.9. These changes are driven by the evolution in the mix of purchasers over the model year, whereby relatively more lower-income (and so more price sensitive) households delay their purchase decisions until later in the year.

Overall, these estimated own-price elasticities are smaller in absolute value compared to the results reported in BLP and Goldberg (1995).[17] It is not straightforward however, 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.

The most comparable article is Copeland, Dunn, and Hall (2011), who use similar transaction-based price data, but employ a static model and work at the quarterly frequency. They find average own-price elasticities around -3.5, which is about twice as large as the estimate found in this article. Most likely, a driver between these estimates is the more expansive definition of the new vehicle market in Copeland, Dunn, and Hall (2011). Specifically, that article considers sales of new vehicles beyond the first 12 months of the model year and exploits the fact that

[17] BLP 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 4: Own-Price Elasticities by Market Segment (absolute value)

Market Segment	Number of models in an average month	Quartiles 1/4	1/2	3/4
CUV	27	1.55	1.70	1.89
Large car	9	1.55	1.65	1.85
Luxury car	54	1.51	1.70	1.90
Midsize car	33	1.49	1.65	1.94
Pickup	16	1.55	1.73	1.94
Compact car	28	1.44	1.75	2.26
SUV	35	1.50	1.67	1.84
Van	17	1.50	1.68	1.87
All	219	1.51	1.69	1.92

Note: CUV is crossover utility vehicle and SUV is sport utility vehicle.

Table 5: Own-Price Elasticities by Month (absolute value)

Month in the Model Year	Quartiles 1/4	1/2	3/4
1	1.43	1.65	1.93
2	1.44	1.65	1.94
3	1.45	1.64	1.95
4	1.45	1.64	1.97
5	1.46	1.64	1.98
6	1.45	1.63	1.99
7	1.55	1.71	1.89
8	1.55	1.71	1.89
9	1.56	1.72	1.90
10	1.56	1.72	1.89
11	1.57	1.72	1.90
12	1.58	1.73	1.90

Note: First month of the model year is September.

automakers simultaneously sell 2 vintages of the same model (where the older vintage sells for a discount). Extrapolating the results from table 5, we expect households which purchase after the twelfth month of the model year to be quite sensitive to price. The exclusion of these price sensitive households, then, may be a main reason why this article finds lower own-price elasticities than Copeland, Dunn, and Hall (2011).

To calculate cross-price elasticities, I considered the case where there is a foreseen 1 percent price increase of a model j in month t. This allows me to compute cross-price elasticities for all models offered for sale leading up to month t, in month t, and after month t. The matrix of all cross-price elasticities is enormous, and so I average them along two dimensions. First, I group elasticities by their distance from the month in which the price increase occurred. Given the 12 month model year, this difference ranges from -11 to +11 months, where a difference of -1 means the elasticity is for a vehicle sold in the month before the price increase occurred (i.e., $t-1$). Second, I group vehicles into three mutually exclusive categories. The first category are vehicles that are the same model (i.e. j), but sold in different months. The second category includes all vehicles in the same market segment as model j, not including model j. Vehicles in this second category are typically considered the closest substitutes to a vehicle j (at least in the cross-section). The third category includes all remaining vehicles. I construct this cross-price elasticity matrix for all (model j, month t) pairs. I then take the average of each cell across all matrices and report the values in table 6 (the elasticities are multiplied by 100 for ease of viewing). Given a 1 percent price increase for a typical car, then, the average cross-price elasticity of the same car sold a month earlier is 0.00128. The average cross-price elasticity for models sold the same month and in the same market segment is 0.00093.

It is important to note that as the number of "months since price change" gets larger in absolute value, the average cross-price elasticity is computed from a smaller number of inferred price elasticities. For 11 months since price change, for example, the average cross-price elasticity is based on changes in sales in the last month of the model year (i.e., month 12), given price changes in the first month of the model year. For 9 months since price change, though, the average cross-price elasticity is based on (i) changes in sales in the month 10 given prices change in month 1, (ii) changes in sales in the month 11 given price changes in month 2, and (iii) changes in sales in the month 12 given price changes in month 3.

The main point to draw from table 6 is the significant amount of intertemporal substitution

Table 6: Average cross-price elasticities

	Elasticities × 100		
Months since price change	Same model	Own Mrkt Seg	Other Mrkt Seg
-11	0.030	0.017	0.014
-10	0.032	0.018	0.015
-9	0.033	0.020	0.016
-8	0.034	0.021	0.017
-7	0.034	0.021	0.017
-6	0.035	0.021	0.017
-5	0.064	0.040	0.035
-4	0.085	0.055	0.048
-3	0.105	0.069	0.061
-2	0.118	0.079	0.070
-1	0.128	0.086	0.077
0	—	0.093	0.082
1	0.121	0.084	0.075
2	0.109	0.075	0.067
3	0.091	0.063	0.055
4	0.072	0.049	0.043
5	0.050	0.034	0.030
6	0.025	0.017	0.014
7	0.024	0.016	0.013
8	0.024	0.016	0.013
9	0.021	0.014	0.011
10	0.020	0.013	0.010
11	0.016	0.011	0.008

Note: The average cross-price elasticities have been multiplied by 100 for ease of viewing. "Same model" is the group of vehicles that are the same model as the vehicle which experienced a price increase, but are sold in different months. "Own Mrkt Seg" is the group of vehicles that are in the same market segment as the vehicle which experienced a price increase (excluding vehicles in the "Same model" group). "Other Mrkt Seg" is the group of all remaining vehicles.

Table 7: Distribution of switching households (percent)

	Total	Months since the price change (absolute value)										
		1	2	3	4	5	6	7	8	9	10	11
intertemporal	73.0	20.8	17.3	13.0	9.0	5.6	2.4	1.9	1.4	0.9	0.5	0.2
cross-sectional	12.1											
outside option	14.9											

Note: "intertemporal" are households who changed the timing of their purchase in response to a price increase, "cross-sectional" are those who switched to a different vehicle within the same month, and "outside option" are households who switched to the outside option. Numbers are the percent of all switching households.

that occurs in response to a price change. Cross-sectional substitution (i.e. the row where "Months since price change" equals 0), although not insignificant, does not stand out as the main choice for switching consumers. Rather, the model predicts that switching consumers, as a group, are most likely to change the timing of their purchase. Not surprisingly, the same model sold in neighboring months is a close substitute; the cross-price elasticities associated with the same vehicle in neighboring months are more than 30 percent higher than the average elasticity of competing vehicles sold in the same month. Consequently, temporary price cuts for a model will significantly cannibalize its sales in surrounding months.

To gain a different perspective on the amount of intertemporal substitution, I focus on consumers who switch their purchase decision when faced with a price increase. For each model j at time t, I compute the absolute value of the number of months by which a switching household changes the timing of their new vehicle purchase. Summing across every (model j, month t) pair, I can compute where, on average, households who switch end up. Overall, 73.0 percent of switching households change the timing of their purchase, with 12.1 percent continuing to buy in the month and 14.9 percent deciding not to buy a new vehicle at any point in the model year (see table 7). The model predicts, then, that intertemporal substitution dominates the cross-sectional movement, with more switching households changing the timing of their purchases by 1, 2, or 3 months than purchasing a different vehicle within the same month. The largest group of switchers, 20.8 percent of all switching households, change the timing of their purchase by 1 month. This is driven, in part, by consumers deciding to buy the same model a month before or after the temporary price change. Although 73.0 percent of switching households change the timing of their purchase, most of these switching households do not dramatically change the timing of their purchase. Indeed, 70 percent of intertemporal switchers purchase a vehicle within 3 months of their original purchase.

The articles of Erdem, Imai, and Keane (2003), Hendel and Nevo (2006a), Hendel and Nevo (2006b), and Gowrisankaran and Rysman (2012) emphasize that dynamic models, unlike their static counterparts, can properly compute short and long-run price elasticities. In particular, for markets where demand dynamics are important, static models will over-estimate long-run price elasticities by not accounting for consumers' ability to time their purchases. This result is also true here. The above price elasticity results emphasize that in reaction to price changes, a significant number of consumers will change the timing of their purchase, but not the vehicle model. When looking at long-run effects then, a static version of the model presented here would over-estimate price elasticities for new vehicle purchases.

Robustness

At this point, it is useful to consider the model's robustness, with a focus on its elasticity estimates. To this end, I estimate two alternative specifications.

The first alternative is a robustness result to test the strong assumption that consumers do not substitute across model years. In particular, if consumers are substituting across model years to a significant degree, then the model's parameter estimates may be biased. My focus is on the distaste-for-price parameter, α, because this is both the main driver behind consumers' decisions of when to buy in the model year, and the key parameter which drives the elasticity estimates. This substitution opportunity seems to be particularly relevant for consumers who are looking at new vehicles in months 10-12 of the model year, right before the new model year vehicles are introduced. Consequently, to test whether shutting off substitution across model years is biasing the parameter estimates, I estimate the model only on the first 9 months of the model year. Under this alternative specification, α is estimated to be 2.32 with a standard error of 0.34, a value which is not statistically different from 2.72 (the estimate in the benchmark case). This robustness check, then, partially assuages concerns that not modeling substitution across model years may bias the estimate of α.

The second alternative gauges the importance of including a half-year dummy as random coefficient. In this specification, we replace the half-year dummy with a random coefficient on height. This second specification does a worse job fitting the data relative to the benchmark

Table 8: Average cash-back rebate over the model year

		Avg	Months within the Automotive Model Year 1	2	3	4	5	6	7	8	9	10	11	12
2006 model	$	893	523	657	799	845	887	871	922	989	1,004	1,054	960	1,103
year	%	3.2	1.9	2.3	2.8	3.0	3.2	3.1	3.3	3.6	3.6	3.8	3.5	3.9
All model	$	835	483	555	648	719	800	833	886	910	965	993	1,016	1,085
years	%	3.3	1.9	2.1	2.5	2.8	3.1	3.3	3.5	3.6	3.8	3.9	4.0	4.3

Note: The first month of the model year is September. "%" stands for percent and is equal to 100 times the ratio of the average cash-back rebate to the average price without rebates."Avg" is average.

case, especially with regard to the income moments. Without the half-year dummy, the model has a smaller decline in income over the model year and a larger estimate of α (3.23 with a standard error of 0.11). This higher estimate of α implies an average own-price elasticity which is more than 10 percent larger in absolute value compared to the benchmark case (1.9 versus 1.7). The inclusion of the half-year dummy, then, is of consequence in that this variable improves the model's fit of the data and lowers the estimate of consumers' distaste-for-price.

6 Sales and revenue effects of incentives

Although automakers influence new vehicle retail prices in a number of ways, cash-back incentives provide a direct channel between manufacturers and consumers.[18] Automakers, especially those headquartered in the U.S., aggressively used cash-back rebates over the 2000-2007 model years; the average cash-back rebate received by consumers over this period was $835, or 3.3 percent of the price excluding rebates. A striking feature of the use of cash-back rebates is their increased size over the model year (see table 8). Cash-back rebates roughly double, in levels and as a percent of price, from the first to twelfth month of the model year. By contributing to the decline in price over the model year, and so increasing consumers' option value to waiting, cash-back rebates influence the distribution of sales over the model year.

[18] Busse, Silva-Risso, and Zettelmeyer (2006) empirical examine how various types of pricing promotions impact the price negotiated between dealers and consumers. For consumer-directed promotions, such as cash-back rebates, they find that 70 to 90 percent of the rebate is captured by consumers.

Table 9: 2006 Model Year Sales and Revenue with and without cash-back rebates

Month	Unit Sales (thousands)				Revenue ($ million)			
	CF	Data	diff.	cum. diff.	CF	Data	diff.	cum. diff.
1	561	538	23	23	15,808	15,008	800	800
2	745	726	19	42	20,778	19,853	925	1,725
3	864	863	1	43	24,440	23,647	793	2,518
4	1,230	1,251	-21	22	35,476	34,607	870	3,388
5	994	1,008	-14	8	27,293	26,561	732	4,120
6	1,124	1,146	-22	-14	30,992	30,244	748	4,868
7	1,386	1,396	-10	-24	38,296	36,838	1,458	6,326
8	1,279	1,294	-15	-39	34,565	33,294	1,272	7,598
9	1,301	1,316	-15	-54	34,691	33,406	1,285	8,883
10	1,297	1,314	-17	-71	34,585	33,327	1,258	10,141
11	1,274	1,268	6	-65	33,719	32,259	1,460	11,601
12	1,118	1,125	-7	-72	29,842	28,686	1,156	12,757

Note: "CF" stands for counterfactual where prices exclude cash-back rebates, "diff." stands for difference between the counterfactual and data, and "cum. diff." stands for cumulative difference.

To measure the intertemporal impact of cash-back rebates on consumers' purchasing behavior, I use the dynamic demand model to predict sales over the model year without these rebates. This comparative static exercise isolates the role of cash-back rebates and their effect on the timing of both sales and revenues earned. This counterfactual is done for the 2006 model year, in which the use of cash-back rebates closely resembles the average over the whole sample period (see table 8).

The aggregate response by consumers to prices without cash rebates is laid out in table 9. Overall, 72 thousand fewer vehicles are sold and $12.8 billion more revenue is earned. The absence of cash-back rebates drives up the sales-weighted average price of a vehicle by 4.2 percent, from $26,255 to $27,365.[19]

The model also predicts a significant dynamic shift, a large movement of consumers from the second half of the model year to the first half. Without incentives, 43 thousand more units are sold within the first three months of the model year and $2.5 billion more revenue is generated. This dynamic response by consumers reflects the decline in the option value

[19]The result whereby revenues increase when prices are raised is consistent with the article's result that demand is elastic, because in the counterfactual all vehicle prices are being raised simultaneously.

of waiting. Without cash-rebates, new vehicle prices decline at a less rapid pace over the model year, lowering the gains from delaying the purchase of a vehicle. On aggregate, then, consumers who purchase in the counterfactual pay higher prices for their vehicles for two reasons. First, without cash-back rebates the prices of all vehicles in all months are higher. Second, because some consumers move up the timing of their purchases in the model year, they pay even higher prices.

This counterfactual illustrates how the automakers' use of cash-back rebates has a substantial dynamic effect on sales by influencing the timing of the household's purchasing decision. It reinforces the importance of intertemporal substitution in this market and, consequently, the importance to automakers in how they set prices over the model year.

7 Conclusion

In this article I estimate a dynamic demand model for light motor vehicles. I find that consumers are price sensitive and willing to change the timing of their new vehicle purchases. Given an expected and temporary price increase, consumers who change their purchase decision are much more likely to change the timing of their purchase than not. That said, consumers do not often radically change when they purchase within the model. I find that on average, more than 60 percent of consumers who switch their purchase decision in response to an expected and temporary price increase, change their timing by 3 months or less.

I then use the dynamic demand model to measure the impact of cash-back rebates on sales over the model year. Because these rebates grow larger over the model year, they influence the timing of all household's purchase decisions. To measure the intertemporal effect of cash-back rebates, I consider the counterfactual where consumers are no longer offered any of these rebates. The model predicts there are 72 thousand less sales, but $12.8 billion more revenue earned. Further, the model predicts a substantial shift in the timing of sales. Because prices now decline by less over the model year, the option value to waiting has decreased. Consequently, a substantial number of households shift their purchase decisions to earlier in the year.

Incorporating dynamics into the demand system for motor vehicles is important to answer a multitude of other questions. In another article 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 help us better understand the connection between primary and secondary markets and to what degree automakers' profitability is impacted by the market for used cars.

References

ACKERBERG, D., AND M. RYSMAN : "Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects," *Rand Journal of Economics*, Vol. 36 (2005), pp. 771–788.

AIZCORBE, A., B. BRIDGMAN, AND J. NALEWAIK : "Heterogeneous Car Buyers: A stylized fact," *Economic Letters*, Vol. 109 (2010), pp. 50–53.

ATTANASIO, O. : "Consumer Durables and Inertial Behavior: Estimation and Aggregation of (S,s) Rules for Automobile Purchases," *Review of Economic Studies*, Vol. 67 (2000), pp. 667–696.

BERRY, S., J. LEVINSOHN, AND A. PAKES : "Automobile Prices in Market Equilibrium," *Econometrica*, Vol. 63 (1995), pp. 841–890.

BERRY, S., J. LEVINSOHN, AND A. PAKES : "Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy," *American Economic Review*, Vol. 89 (1999), pp. 400–430.

BUSSE, M., C. KNITTEL, AND F. ZETTELMEYER : "Are Consumers Myopic? Evidence from New and Used Car Purchases," *American Economic Review*, Vol. 103 (2013), pp. 220–256.

BUSSE, M., J. SILVA-RISSO, AND F. ZETTELMEYER : "$1000 Cash Back: The Pass-Through of Auto Manufacturer Promotions," *American Economic Review*, Vol. 96 (2006), pp. 1253–70.

BUSSE, M., D. SIMESTER, AND F. ZETTELMEYER : "The Best Price You'll Ever Get: The 2005 Employee Discount Pricing Promotions in the US Automobile Industry," *Marketing Science*, Vol. 29 (2010), pp. 268–290.

CHEN, J., S. ESTEBAN, AND M. SHUM : "When Do Secondary Markets Harm Firms?," *American Economic Review*, Vol. 103 (2013), pp. 2911–2934.

COPELAND, A., W. DUNN, AND G. HALL : "Inventories and the Automobile Market," *Rand Journal of Economics*, Vol. 42 (2011), pp. 121–149.

COPELAND, A., AND G. HALL : "The Response of Prices, Sales, and Output to Temporary Changes in Demand," *Journal of Applied Econometrics*, Vol. 26 (2011), pp. 232–269.

CORRADO, C., W. DUNN, AND M. OTOO : "Incentives and Prices for Motor Vehicles: What Has Been Happening in Recent Years?," Finance and Economics Discussion Series 2006-09, Board of Governors of the Federal Reserve System, 2006.

ERDEM, T., S. IMAI, AND M. KEANE : "A Model of Consumer Brand and Quality Choice Dynamics Under Price Uncertainty," *Quantitative Marketing and Economics*, Vol. 1 (2003), pp. 5–64.

GOLDBERG, P. K. : "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," *Econometrica*, Vol. 63 (1995), pp. 891–951.

GOLDBERG, P. K., AND F. VERBOVEN : "The Evolution of Price Dispersion in the European Car Market," *Review of Economic Studies*, Vol. 67 (2001), pp. 811–848.

GOWRISANKARAN, G., AND M. RYSMAN : "Dynamics of Consumer Demand for New Durable Goods," *Journal of Political Economy*, Vol. 120 (2012), pp. 1173–1219.

GRAMLICH, J. : "Gas Prices, Fuel Efficiency, and Endogenous Product Choice in the U.S. Automobile Industry," Yale University manuscript, 2010.

HALL, G. : "Non-Convex Costs and Capital Utilization: A Study of Production Scheduling at Automobile Assembly Plants," *Journal of Monetary Economics*, Vol. 45 (2000), pp. 681–716.

HENDEL, I., AND A. NEVO : "Sales and Consumer Inventory," *Rand Journal of Economics*, Vol. 37 (2006a), pp. 543–561.

HENDEL, I., AND A. NEVO : "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, Vol. 74 (2006b), pp. 1637–1673.

HENDEL, I., AND A. NEVO : "Intertemporal Price Discrimination in Storable Goods Markets," *American Economic Review*, Vol. 103 (2013), pp. 2722–2751.

KAHN, J. : "Inventories and the Volatility of Production," *American Economic Review*, Vol. 77 (1987), pp. 667–79.

KAHN, J. : "Why is Production More Volatile than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory-Holding," *Quarterly Journal of Economics*, Vol. 109 (1992), pp. 565–92.

LANGER, A., AND N. MILLER : "Automakers' Short-Run Responses to Changing Gasoline Prices," *Review of Economics and Statistics*, Vol. 95 (2013), pp. 1198–1211.

MAGNAC, T., AND D. THESMAR : "Identifying Dynamic Discrete Decision Processes," *Econometrica*, Vol. 70 (2002), pp. 801–816.

PASHIGIAN, P., B. BOWEN, AND E. GOULD : "Fashion, Styling, and the Within Season Decline in Automobile Prices," *Journal of Law and Economics*, Vol. 38 (1995), pp. 281–310.

PESENDORFER, M. : "Retail Sales: A Study of Pricing Behavior in Supermarkets," *Journal of Business*, Vol. 75 (2002), pp. 33–66.

PETRIN, A. : "Quantifying the Benefits on New Products: The Case of the Minivan," *Journal of Political Economy*, Vol. 110 (2002), pp. 705–729.

RAMEY, V. A., AND D. J. VINE : "Declining Volatility in the U.S. Automobile Industry," *American Economic Review*, Vol. 96 (2006), pp. 1876–1889.

SCHIRALDI, P. : "Automobile Replacement: a Dynamic Structural Approach," *Rand Journal of Economics*, Vol. 42 (2011), pp. 266–291.

SHUM, M., AND S. ESTEBAN : "Durable Goods Oligopoly with Secondary Markets: the Case of Automobiles," *Rand Journal of Economics*, Vol. 38 (2007), pp. 332–354.

SLADE, M. E. : "Optimal Pricing with Costly Adjustment: Evidence from Retail-Grocery Prices," *Review of Economic Studies*, Vol. 65 (1998), pp. 87–107.

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

Appendix A

In this appendix, I detail the regressions I ran to determine whether old model-year prices have an effect on new model-year sales. These regressions were performed in the context of assessing the modeling assumption of no substitution across model-years. As detailed below, the regression results find little effect of old model-year prices on new model-year sales, consistent with the claim that new and old model year vehicles of the same model are poor substitutes.

Because prices and sales have trends over the model, I take a two-stage approach to measuring if old model-year prices impact new model-year sales. In the first stage I 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 month and the model year, for each model in the sample.[20] Let $(\hat{s}^y_{jt}, \hat{p}^y_{jt})$ denote the resulting sales and price residuals, respectively, for a model j of model year y in month t.

In the second stage, I explore whether there is a correlation between the change of $\hat{s}^y$ and the change of $\hat{p}^{y-1}$, controlling for the change in $\hat{p}^y$. I ran the following regression for $n = \{1,3,5\}$ and for the model year pairs $\{(2007,2006);(2005,2004);(2002,2001)\}$,

$$\hat{s}^y_{jt} - \hat{s}^y_{j,t-n} = \alpha_0 + \alpha_1(\hat{p}^y_{jt} - \hat{p}^y_{j,t-n}) + \alpha_2(\hat{p}^{y-1}_{jt} - \hat{p}^{y-1}_{j,t-n}). \tag{9}$$

For the 9 regressions I ran, α_2 was statistically insignificant for 7 of them. In contrast, α_1 was negative and statistically significant for 8 of the 9 regressions (see table 10).

Taken all together, these reduced form results suggest that the typical consumer does not consider new and old model-year vehicles of the same model to be close substitutes. The model's assumption of no substitution between model-years, then, seems reasonable.

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

Table 10: Lagged relationship across model years

	Dependent variable		
	$\hat{s}^{2007}_{jt} - \hat{s}^{2007}_{j,t-1}$	$\hat{s}^{2007}_{jt} - \hat{s}^{2007}_{j,t-3}$	$\hat{s}^{2007}_{jt} - \hat{s}^{2007}_{j,t-5}$
constant	0.065*	0.048*	0.074*
$\hat{p}^{2007}_{jt} - \hat{p}^{2007}_{j,t-1}$	-1.148*		
$\hat{p}^{2007}_{jt} - \hat{p}^{2007}_{j,t-3}$		-2.676*	
$\hat{p}^{2007}_{jt} - \hat{p}^{2007}_{j,t-5}$			-2.926*
$\hat{p}^{2006}_{jt} - \hat{p}^{2006}_{j,t-1}$	-0.255		
$\hat{p}^{2006}_{jt} - \hat{p}^{2006}_{j,t-3}$		-0.720*	
$\hat{p}^{2006}_{jt} - \hat{p}^{2006}_{j,t-5}$			-1.213*
obs.	1,675	1,323	960

	$\hat{s}^{2005}_{jt} - \hat{s}^{2005}_{j,t-1}$	$\hat{s}^{2005}_{jt} - \hat{s}^{2005}_{j,t-3}$	$\hat{s}^{2005}_{jt} - \hat{s}^{2005}_{j,t-5}$
constant	0.080*	0.082*	0.041*
$\hat{p}^{2005}_{jt} - \hat{p}^{2005}_{j,t-1}$	-2.081*		
$\hat{p}^{2005}_{jt} - \hat{p}^{2005}_{j,t-3}$		-2.496*	
$\hat{p}^{2005}_{jt} - \hat{p}^{2005}_{j,t-5}$			-1.870*
$\hat{p}^{2004}_{jt} - \hat{p}^{2004}_{j,t-1}$	-0.422		
$\hat{p}^{2004}_{jt} - \hat{p}^{2004}_{j,t-3}$		-0.329	
$\hat{p}^{2004}_{jt} - \hat{p}^{2004}_{j,t-5}$			-0.084
obs.	1,770	1,376	987

$\hat{s}^{2002}_{jt} - \hat{s}^{2002}_{j,t-1}$	$\hat{s}^{2002}_{jt} - \hat{s}^{2002}_{j,t-3}$	$\hat{s}^{2002}_{jt} - \hat{s}^{2002}_{j,t-5}$	
constant	0.041*	-0.054*	-0.128*
$\hat{p}^{2002}_{jt} - \hat{p}^{2002}_{j,t-1}$	-1.005*		
$\hat{p}^{2002}_{jt} - \hat{p}^{2002}_{j,t-3}$		-1.325*	
$\hat{p}^{2002}_{jt} - \hat{p}^{2002}_{j,t-5}$			-1.162
$\hat{p}^{2001}_{jt} - \hat{p}^{2001}_{j,t-1}$	0.116		
$\hat{p}^{2001}_{jt} - \hat{p}^{2001}_{j,t-3}$		0.488	
$\hat{p}^{2001}_{jt} - \hat{p}^{2001}_{j,t-5}$			-0.254
obs.	1,537	1,196	851

Notes: Let $(\hat{s}_{jt}, \hat{p}_{jt})$ 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 for a model j. $*$ Significant at the 5 percent level.