1-2010

Drivers of Finished Goods Inventory in the U.S. Automobile Industry

Gerard. P. Cachon
University of Pennsylvania

Marcelo Olivares

Recommended Citation
Cachon, G. P., & Olivares, M. (2010). Drivers of Finished Goods Inventory in the U.S. Automobile Industry. *Management Science, 56* (1), 202-216. http://dx.doi.org/10.1287/mnsc.1090.1095

Drivers of Finished Goods Inventory in the U.S. Automobile Industry

Abstract
Automobile manufacturers in the U.S. supply chain exhibit significant differences in their days of supply of finished vehicles (average inventory divided by average daily sales rate). For example, from 1995 to 2004, Toyota consistently carried approximately 30 fewer days of supply than General Motors. This suggests that Toyota's well-documented advantage in manufacturing efficiency, product design, and upstream supply chain management extends to their finished-goods inventory in their downstream supply chain from their assembly plants to their dealerships. Our objective in this research is to measure for this industry the effect of several factors on inventory holdings. We find that two factors, the number of dealerships in a manufacturer's distribution network and a manufacturer's production flexibility, explain essentially all of the difference in finished-goods inventory between Toyota and three other manufacturers: Chrysler, Ford, and General Motors.

Keywords
empirical, supply chain management, distribution, product variety, inventory theory, manufacturing flexibility

Disciplines
Marketing | Operations and Supply Chain Management

Columbia Business School

Decision, Risk & Operations
Working Papers Series

Drivers of finished goods inventory in the U.S. automobile industry

G. P. Cachon and M. Olivares

October 2006; revised January 2008
DRO-2009-03

Drivers of finished goods inventory in the U.S. automobile industry*

Gérard P. Cachon · Marcelo Olivares

The Wharton School, University of Pennsylvania, Philadelphia PA, 19104
cachon@wharton.upenn.edu
opim.wharton.upenn.edu/~cachon

Columbia Business School, Columbia University, New York NY,
molivares@columbia.edu
www.columbia.edu/~mo2338/index.htm

October 24, 2006; Revised January 31, 2008

Abstract

Automobile manufacturers in the U.S. supply chain exhibit significant differences in their days-of-supply of finished vehicles (average inventory divided by average daily sales rate). For example, from 1995 to 2004, Toyota consistently carried approximately 30 fewer days-of-supply than General Motors. This suggests that Toyota's well documented advantage in manufacturing efficiency, product design and upstream supply chain management extends to their finished-goods inventory in their downstream supply chain from their assembly plants to their dealerships. Our objective in this research is to measure for this industry the effect of several factors on inventory holdings. We find that two factors, the number of dealerships in a manufacturer's distribution network and a manufacturer's production flexibility, explain essentially all of the difference in finished goods inventory between Toyota and three other makes, Chrysler, Ford and General Motors.

Keywords: Empirical, supply chain management, distribution, product variety, inventory theory, manufacturing flexibility.

*The authors thank the Associate Editor and three reviewers for their comments. Financial support for this research was provided by the Fishman-Davidson Center, The Wharton School, University of Pennsylvania.

1. Introduction

The auto industry is clearly important to the overall world economy and it has been a source of many innovations in manufacturing technology (e.g., the assembly line, just-in-time inventory, kan-ban, etc.) and product design. As a result, it has been the subject of numerous empirical studies. However, most of these studies have been centered on analyzing the production and procurement processes (e.g., Lieberman et al. (1990) and Lieberman and Asaba (1997)) or the new product development process (e.g., Clark and Fujimoto (1989)). Little attention has been placed on the management of the finished goods from the assembly plant down to the consumer, which is the focus of this paper.

Figure 1 displays times series of the days-of-supply (end of month inventory divided by the average daily sales rate on the following two months) for three auto manufacturers– Ford, General Motors (GM) and Toyota– between 1995 and 2004. This measure of inventory performance includes all finished goods inventory destined for sale in the U.S. market and physically in North America: inventory on factory lots, at ports of entry, in-transit to dealerships and at dealerships. The figure reveals striking differences among the different makes. Although on average the makes hold about 60 days-of-supply (which is often suggested in the trade press as the "ideal" inventory level, Harris (2004)), Toyota consistently holds less than that benchmark while GM and Ford hold more than that benchmark in the majority of the sample. Furthermore, none of the companies exhibit a trend in inventory during this time period, which suggests that these differences are persistent[1]. Our objective in this study is to measure the effect of several factors that could explain the differences in inventory observed in the industry.

Based on analytical models and empirical studies in the operations management literature, we identify numerous factors that could influence a firm's optimal inventory decision. These factors can be roughly group into four categories: demand fragmentation, sales characteristics, production characteristics and competition. *Demand fragmentation* refers to the allocation of demand across different products (i.e., vehicle models), or across different options of a given product or across different geographic locations (e.g., dealerships). Each of these forms of fragmentation can lead to more variable demand and therefore more

[1] We regressed days-of-supply on a linear time trend and monthly dummies assuming AR(1) errors. Our analysis suggests that only five of the fifteen manufacturers exhibit a trend. Among the six major manufacturers, only Nissan exhibits a (negative) trend. Porsche and Isuzu are the only manufacturers that have trends in nominal inventories (both positive).

inventory. For *sales characteristics* we focus on sales trends and seasonality. *Production characteristics* refer to a firm's manufacturing capabilities and production schedule. For example, a more flexible plant can adjust its production more readily and therefore can better match its production to its sales. Hence, a more flexible plant could enable a firm to hold lower safety stocks. Furthermore, holding a plant's production flexibility constant, inventory should increase when the plant is required to produce a greater variety of products, due to switching times between products. *Competition* refers to the amount of competition a firm faces for its products. We conjecture this can influence a firm's inventory in at least two ways: (1) competition should reduce a product's margin, which leads to lower inventory; and/or (2) competition gives consumers more choices, which leads to higher inventory - when a consumer has choices it is important to have in stock a product that closely matches the consumer's preference, otherwise the consumer is more likely to substitute to a competitor's product. Although theory enables us to identify these various factors, an empirical study is needed to evaluate their relative importance (at least for our focus industry and market, U.S. autos).

The next section reviews the related literature. Section 3 gives a brief introduction to the industry and section 4 describes the data used. Section 5 describes the factors included in our econometric model and the estimation methods. Section 6 provides our estimation results and sensitivity analysis. Section 7 discusses our main conclusions and future related work.

2. Literature Review

Most studies of operational performance in the auto industry have focused within the assembly plant or on the product design process rather than finished goods in the downstream supply chain. For example, Fisher and Ittner (1999) measure the effect of product variety on work-in-process inventory using archival data from automotive plants of a single company. MacDuffie et al. (1996) analyze the impact of product variety on manufacturing productivity and consumer-perceived quality using data from 70 auto assembly plants. Lieberman et al. (1990) analyze drivers of productivity growth across firms in the auto industry, which includes labor, capital and total factor productivity. Lieberman and Demeester (1999) demonstrate that reductions in work-in-progress inventory can lead to productivity gains, which is a causal relationship that is econometrically challenging to identify due to the feedback between the two variables. Lieberman and Asaba (1997) report interesting differences regarding inventory performance across the supply chains of Japanese and U.S. auto manufacturers, but they exclude finished goods inventory from the analysis. Clark and Fujimoto (1989) study the effect of several product and project characteristics and organizational capabilities on new product development leadtimes. Bresnahan and Ramey (1994) and Hall (2000) provide evidence of significant adjustment costs in the production rate at auto plants, leading manufacturers to have intermittent plant closings to match supply with demand. Goyal et al. (2006) study factors that influence the adoption of flexible production technology by U.S. auto manufacturers. We add to this stream of research by linking other factors associated with production and scheduling that are associated with finished-goods inventory.

Several papers explore inventory at the industry level with a focus on either the long run trend in inventory (e.g., Wu et al. (2005) and Rajagopalan and Malhotra (2001)) or the volatility of production relative to sales (e.g., Cachon et al. (2007)) - we do not consider either of those issues in our study. There is a growing literature that explores firm level inventory rather than at the product/model level as we do. For example, Gaur et al. (2005), use panel data from quarterly financial reports of retailers to find that inventory turnover is negatively related to a retailer's capital intensity and positively related to the retailer's gross margin and a proxy for sales forecast errors. We focus on finished-goods inventory performance over a larger section of the supply chain (assembly plant down to retailer/dealer) and because we concentrate on one product category (automobiles), we are able to obtained more detailed data on other factors that influence inventory performance. Rumyantsev and Netessine (2007) use aggregate inventory data of public U.S. companies to measure the relationship between demand uncertainty, lead times, gross margins and firm size on inventory levels. We include similar covariates in our study. Hendricks and Singhal (2005), Wu et al. (2005), Lai (2006) and Randall et al. (2006) study the relationship between inventory and firm financial performance measures, but we do not consider such measures (again, because our unit of analysis is the product/model level rather than the company level).

3. The U.S. Automotive Industry

In this section we provide a brief description of some idiosyncratic features of the U.S. auto industry that are important for our study. Six companies account for about 90% of sales in the U.S. auto market: Chrysler, Ford, General Motors (GM), Honda, Nissan and Toyota.[2] More than 90% of U.S. sales for Chrysler, Ford and GM is produced in the U.S., Canada and Mexico. We refer to vehicles produced in (outside) North America as domestic (imported). Toyota and Honda produce about 50% of their U.S. sales domestically, while 65% of Nissan's vehicles are domestic. Some companies, e.g. Hyundai and Porsche, satisfied all of their U.S. sales with imported production during our study period.

There are different levels of aggregation at which one can describe product variety in the auto industry. Each company offers vehicles under several brands or auto *makes*. For example, GM makes include Chevrolet, GMC and Pontiac, among others; Toyota makes are Toyota division (hereafter Toyota), Lexus and Scion. Each auto make produces several auto *models.* Examples of models include the Chevrolet Cavalier, the Toyota Camry and the Ford Explorer. Models can be classified into vehicle *types*, which include cars, sport cars, sport utility vehicles, pickups, minivans, etc. A *platform* is often used to describe commonality among models at the production level. For example, the Harbour Report (2004, pg. 229) defines a platform as the "welded or framed underbody a car is built and rides on" and designates that the Chevrolet Cavalier and the Pontiac Sunfire are built on the same platform. Consumers purchase models with different *options,* which include different body styles, engines, transmission types, safety features (e.g., side airbags, automatic breaking system) and other accessories.

Automobile assembly plants consist of one or more assembly lines that are designed to produce in large scale a particular vehicle specification with a limited range of options. Opening a new assembly plant requires significant capital investment and assembly lines are designed to operate at a particular line rate (vehicles per hour). As a result, in the short-run, a manufacturer's primary option for adjusting production is either to add or to subtract shifts (Bresnahan and Ramey (1994)).

Franchise laws regulate new vehicle sales in the U.S. and all new vehicles must be sold through a network of dedicated franchised dealers. In the U.S. most vehicles are purchased directly from dealership inventory[3]. Furthermore, dealerships do not order inventory like retailers in most other industries, but rather manufacturers implement a push system that allocates inventory to dealerships after production (e.g., Cachon and Lariviere (1999)).

[2] Chrysler merged with Daimler-Benz in 1998, changing its name to Daimler-Chrysler, but we continue to refer to the company throughout as Chrysler.

[3] Marti (2000) reports that only 15-20% of buyers buy custom cars from manufacturers.

Hence, we study the performance of all finished goods inventory in the supply chain from the assembly plant down to the dealership.

4. Data

We collected data, covering the years 1996 through 2004, from three main secondary sources: Automotive News, Ward's Auto and Harbour Report. From Wards Auto, we obtained monthly end-of-the-month inventory and sales by model. Inventory includes all finished automobiles in North America destined for sale in the U.S. market: inventory on factory lots and ports of entry, inventory in-transit to dealerships and inventory at dealerships [4] We also obtained (i) model specifications and list prices for all cars and light-trucks (pickups, vans and SUVs) available by year, (ii) monthly domestic production of each model by plant, and (iii) the platform designations of each model. From Automotive News we obtained data on (i) the number of dealerships by auto make by year, (ii) survey data on gross profits of dealerships by auto make by year, and (iii) model specifications which were used to complete and cross validate the data published by Wards. We also obtained data on plant stoppages from the weekly periodicals of Automotive News.

From The Harbour Report we obtained data on a selection of assembly plants in North America. Several plants have more than one production line and the data is reported separately for each line. In those cases we refer to each production line as a distinct plant. The data include total production, line rate capacity and the number of platforms produced by plant by year.[5] We also have data on the models that were produced at each plant. Harbour includes data for all Chrysler, Ford and GM plants, with the exception of Chrysler's Conner Avenue plant. The Harbour Report does not include plants from BMW, Mercedes, Subaru, Volkswagen and Volvo. The plants in the Harbour Report cover 90% of total domestic production during year 1996 through 2004. Coverage is excellent for Chrysler, Ford and GM but somewhat lower for Toyota and Honda due to the exclusion of some of their plants[6].

In addition to these data, we obtained some economic data, such as the price of gasoline, consumer price indexes, number of households in the U.S. and personal income data. These

[4] Exports are a small fraction of U.S. production and are often shipped as parts, therefore not counted as finished vehicles. GM changed its inventory counting scheme during the study period, reporting dealership inventory only. We included dummy variables to control for this change in our econometric study.

[5] For three plants, the number of platforms is provided for the plant and not for each production line. All our results are robust to the exclusion of these plants.

[6] Tables A1 and A2 in the online appendix describes in more detail the plants included in the The Harbour Report.

were obtained from the Current Population Survey, Energy Information Administration and from the Bureau of Labor Statistics. We collected data from Consumer Reports for our sensitivity analysis.

We excluded some data in our econometric analysis. The Chevrolet Lumina was phased-out in years 2000-2001 and sold only to rental companies, so we chose to exclude it. We also excluded the Chevrolet Metro in 2001 (its last year of production) and the Saturn EV1 (an electric vehicle), both of which had a days-of-supply greater than 600 (more than 20 standard deviations above the mean). We excluded the Ford Excursion in 2000 because its plant utilization was more than 5 standard deviations above the mean. GM Oldsmobile and Suzuki had the largest variation in the number of dealerships during the study period. GM announced the closing of Oldsmobile in 2000 and the last model was produced in 2004 (the number of dealerships was reduced from 2990 to 1337). Suzuki experience the opposite change in its dealership structure - it expanded from 290 dealers in 1995 to 543 in 2005. We chose a conservative approach and excluded from our main results observations from Oldsmobile from 2000-2004 and all Suzuki observations because these dramatic changes in the dealership structure could be correlated with other factors that affect inventory (e.g., such as closing a brand or building a brand).[7] We also excluded full sized vans and pickups from our analysis because models in these segments tend to exhibit huge option variety (e.g. Ford F-Series has an average of 280 options offered per year). As we show later, our estimation requires data from assembly plants, so our sample only includes models produced at plants covered by the Harbour Report.

5. Econometric Specification

This section describes the measures used in our study, our hypotheses and the estimation of our empirical model.

We use i to index vehicle models (hereon models) and t to index calendar years (hereon years). The dependent variable is the log of the average monthly days-of-supply, DS_{it}, of each model in each year, where days-of-supply in a month equals the inventory at the end of the month divided by the average daily sales rate in the following two months. Specifically, for models that were sold in each month of a year,

[7] Section 6.1 shows some results when these makes are included.

$$DS_{it} = \frac{1}{12} \sum_{m=1}^{12} \left(\frac{I_{itm}}{\sum_{k=m+1}^{m+2} S_{itk}} \right),$$

where I_{itm} is end-of-month inventory (in units) in month m and S_{itm} is sales (in units) in month m. (Naturally, months 13 and 14 in year t are actually months 1 and 2 in year $t+1$.) If a model was sold for part of a year, we average the days-of-supply from only those months. Finally, the average days-of-supply does not include the last two months a model is sold. We use a forward looking assessment of the sales rate (two months ahead) because we expect that inventory is held in anticipation of future demand rather than in reaction to past demand, especially when demand exhibits known seasonal patterns. Our results are robust to alternative measures of days-of-supply.[8] A log transformation is consistent with previous studies (e.g., Gaur et al. (2005)), but we report in section 6.1 results without a log transformation.

The independent variables are divided into two groups: measures associated with individual models, denoted by the (column) vector X_{it} and measures attributed to the plant producing a model, $p(i)$, denoted by the (column) vector $W_{p(i)t}$. The third group in our model is an error term, u_{it}, that captures unobserved factors and other random fluctuations affecting DS. Thus, the econometric model is defined as:

$$DS_{it} = \beta X_{it} + \gamma W_{p(i)t} + u_{it} \ , \tag{1}$$

where β and γ are row vector parameters to be estimated. Like DS, all variables in X and W are included with log transformation. We next detail the particular measures included in X and W. Subsequently, we divide u_{it} into additional components. As mentioned in the introduction, our measures can linked to factors from four broad (and not necessarily mutually exclusive) categories: demand fragmentation, sales characteristics, production characteristics and competition.

Three of the factors included in X_{it} are related to various forms of demand fragmentation: $SALES_{it}$, $OPTIONS_{it}$, and $DEALERS_{it}$. $SALES_{it}$ is the average monthly sales (in units) of model i during year t (again, only including months for which the model was sold): as a brand adds models to its assortment it may reduce the annual sales per model as its aggregate sales become fragmented over its wider product offering. $OPTIONS_{it}$ is the number of options

[8] We considered three other methods for evaluating the sales rate in the denominator of the days-of-supply ratio: (1) the average sales in the following month only, (2) the average sales rate in the following three months; and (3) sales in the same month inventory is measured. Our results with these measures were similar, but the estimates were less precise in particular when DS was calculated using the third option.

offered for model i in year t, where an increase in a model's options may be associated with fragmenting its inventory into units that are not perfect substitutes. The definition of these options are relatively standard, so it is possible to make comparisons of option intensity across models and years.[9] Finally, $DEALERS_{it}$ is the number of dealerships in year t of model i's brand: all else equal, an increase in the number of dealerships fragments sales among more physical locations. If there are economies of scale associated with inventory and production, then we expect that DS should be decreasing in $SALES$: models with a higher sales volume may require proportionally less inventory. (For example, it is well known that the EOQ model exhibits economies of scale - doubling demand increases inventory by less than a factor of 2.) Similarly, DS should be increasing in $OPTIONS$ and $DEALERS$.[10] Furthermore, an increase in $DEALERS$ is likely to increase the competition faced by a brand's dealerships, both intra and inter-brand competition. Theory is ambiguous with respect to the impact of competition on DS, but there is some evidence that in isolated markets there is a positive association between the level of competition and DS (see Olivares and Cachon (2007) for details).

In addition to the nominal level of sales, X_{it} includes covariates capturing sales trends. Production capacity can be costly to adjust in the short run, so changes in sales from year to year may lead to deviations from target inventory levels. For example, we expect a model's DS decreases when sales increase from one year to the next as production capacity may lag the sales growth. Alternatively, a model with increasing sales may be a popular product and therefore require less inventory. To explain, when a consumer does not find her most preferred product, she can choose not to purchase anything, thereby causing a "lost sale", or choose to wait for the product to become available (i.e., backorder) or purchase another product in the assortment. A consumer will be more patient with a more popular/well designed product (i.e., more likely to backorder rather than choose the lost-sales option). Therefore, a firm with a popular product can rationally choose to carry less inventory. We cannot directly observe a model's "popularity" (our catch-all phrase for the propensity of consumers to wait to purchase an item rather than substitute), but we suspect it may be associated with several of our measures, as we will discuss.

[9] Kekre and Srinivasan (1990) uses cross-sectional survey data from different industries to estimate the effect of product variety on inventory, but finds no significant impact. Measuring differences in product variety across industries is challenging and could be causing this negative result. By focusing on the auto industry, we are able to use more detailed and objective measures of product variety.

[10] For example, fragmenting demand may increase the coefficient of variation of demand (as assumed by van Ryzin and Mahajan (1999)), thereby requiring more inventory to achieve the same target service level.

Continuing with the issue of sales trends, we include in X_{it} the following two measures:

$$STREND_{it}^{+} = \max\left((SALES_{it} - SALES_{it-1})^{+}, 1\right) \tag{2a}$$

$$STREND_{it}^{-} = \max\left((SALES_{it-1} - SALES_{it})^{+}, 1\right), \tag{2b}$$

where x^{+} denotes $\max(x, 0)$. (These measures are never less than 1, which ensures that we can apply a log transformation to each of them.) We expect DS is decreasing in $STREND^{+}$ (either because of production reasons or because of model popularity) and, naturally, DS should be increasing in $STREND^{-}$. We include two measures to allow for different reactions to sales increases and decreases.[11]

Sales in the auto industry exhibit varying degrees of seasonality, which motivates a production smoothing strategy when it is costly to change the level of production - produce at a reasonably constant level, build up inventory during slow sales periods and draw down inventory during sales peaks. As a result, we expect that DS is increasing in the degree of seasonality - the more seasonal sales are, all else being equal, the more inventory a firm rationally carries. Alternatively, sales seasonality could be a proxy for model popularity. A popular product may not exhibit strong seasonal markdown patterns (i.e., discount routinely offered at the same time of the year) and therefore a popular product may exhibit less seasonality. Consequently, DS may increase with seasonality because a seasonal product is unpopular, thereby requiring more inventory to avoid lost sales.

To measure seasonality, with each sales-time series, we fit a regression with model-specific monthly dummies, denoted d_{im}, $m \in \{1, ..., 12\}$. Our seasonality measure for model i is

$$SEASON_i = \sqrt{V(d_{im})}/E(S_{itm}), \tag{3}$$

where $V(\cdot)$ denotes the sample variance, and $E(\cdot)$ the sample mean.

As we have already discussed, a model's popularity may be associated with a sales trend or sales seasonality. We believe there can be other factors associated with a model's popularity. For example, the number of models within a segment may be a measure of product substitution. As more models are added to a segment, they become closer to each other

[11] Our sales trend measures begin in 1996 because our sales data begins in 1995. Some new models were introduced during our study period. Usually, sales of a new model start in the second half of the year previous to the model-year of introduction. For example, the Cadillac Escalade was launched in model-year 1999, but sales for this model started on October of 1998. For this model, $STREND$ is calculated for 1999 as the difference in average monthly sales between 1999 and 1998. Similar calculations were used for the other new models. Excluding models in their year of introduction does not change our main results.

in a product attribute space. Consequently, consumers will be less likely to wait if their preferred model is out of stock because their second choice becomes a more attractive alternative. Therefore, we include in X_{it} the number of models in the same segment as model i in year t, $NMKT_{it}$, as a measure of product substitution and competition. We expect that DS is increasing in $NMKT$.

Next, model popularity may be correlated with a model's cost markup (a model's gross margin divided by its marginal cost) - a firm can rationally support a higher cost markup with a popular product. Alternatively, a firm may carry more inventory of a product with a high markup because the consequence of a lost sale is greater - a stockout can be costlier with a high gross margin product than a low gross margin product. Hence, the association between cost markup and inventory is ambiguous– the direct effect on the cost of lost sales tends to increase inventory as markups increase, but the indirect effect of model popularity tends to decrease inventories. We are not able to observe $COSTMK_{it}$, the cost markups for model i in year t, so, following Berry et al. (1995), we estimate the cost markups for each model using a structural model of oligopoly price competition in a differentiated product market. In short, this methodology estimates the cross-price elasticities among all products offered during a year, and computes equilibrium markups based on competitive pricing under the estimated demand system. Details of the implementation are described in the online appendix[12].

Finally, we include in X_{it} a measure of production flexibility. We do not observe production flexibility directly, so we seek to observe the application of flexibility. In particular, if a model is manufactured in a flexible plant, then we conjecture the plant is able to produce in small batches, switch production between models without substantial downtime periods and/or possibly increase or decrease production by adding or subtracting shifts and/or overtime. As a result, production in a flexible plant should track sales more closely relative to an inflexible plant. Therefore, we proxy production flexibility by the average absolute difference between production and sales, normalized by sales volume,

$$PS_{it} = \frac{E\left(|P_{itm} - S_{itm}|\right)}{E\left(S_{itm}\right)} = \frac{E\left(|I_{itm} - I_{itm-1}|\right)}{E\left(S_{itm}\right)}, \tag{4}$$

where P is the production series, and the equation above follows from inventory balance (i.e., the change in inventory equals the difference between production and sales).[13] Consistent with our assertion that inventory volatility (i.e., a higher PS measure) proxies for

[12] We use part of the computer code developed by Nevo (2000) to do the estimation.

[13] We use the inventory series rather than the production and sales series because some of North American production is not sold in the U.S., especially for plants in Mexico. As a result, the production series (for all

production flexibility, we found a positive correlation between the PS of models produced at the same plant (0.32). To explain further, consider the typical saw-tooth inventory pattern implied by a (Q, r) inventory policy (e.g. Nahmias (2005), pg. 251). Inventory depends both on the level of safety stock (the amount of inventory at the inventory troughs) as well as on the size of the batches. An inflexible production process produces in large batches and therefore exhibits more inventory volatility. Furthermore, there need not be a mechanical, or one-to-one, relationship between the average inventory level (DS) and the amount of inventory volatility (PS)– two products can have the same average inventory level but different inventory volatilities, or two products can have the same inventory volatility (batch size) but different average inventory levels (because they carry different safety stock levels). Nevertheless, we expect that PS has a positive effect on DS – all else being equal, a firm with lower inventory volatility carries less inventory.

Now consider $W_{p(i)t}$, which includes characteristics of the plants that produce model i. To account for the time to switch between producing different models, $NPLATF_{p(i),t}$ is the number of platforms (as defined by the Harbour Report) produced at plant $p(i)$ in year t. For models that were produced at more than one plant during the same year, $p(i)$ denotes a weighted average plant, calculated with production quantities as weights. We expect $NPLATF$ to have a positive effect on DS, due to production switching times.[14] A measure of capacity utilization, $UTIL$, is also included in $W_{p(i)t}$. We calculated $UTIL$ assuming a constant per hour production rate of the plant during the year (using Harbour's line rate measure), three 8-hour shifts and 365 days per year. Theory is ambiguous on the effect of $UTIL$ on inventory performance.[15]

The third group in (1) is the error term, which we decompose into different random components:

$$u_{it} = \delta_i + \omega_{p(i)} + \tau_t + \epsilon_{it}^m + \epsilon_{p(i)t}^w \; . \tag{5}$$

The random components δ_i and $\omega_{p(i)}$ represent time-invariant unobserved factors related to

of North America) and the sales series may not balance with the inventory series (for just the U.S. market). Because we are studying the U.S. supply chain, we prefer to base our measure of flexibility on the changes in U.S. inventory.

[14] For example, in an economic lot scheduling problem with cyclic schedules, adding platforms to a production process requires an increase in the production batches, which leads to higher inventory.

[15] In a make-to-stock queuing model an increase in utilization increases a product's lead time, which can increase the inventory needed to maintain a target service level. This suggests a positive relationship between $UTIL$ and DS. However, consider a cyclic production schedule with multiple products and switching times between products. If there is a minimum production quantity (e.g., one shift), then there can be a negative relationship between $UTIL$ and DS.

model i and plant $p(i)$ where the model is produced, respectively. The term τ_t represents time shocks that affect inventory performance industry-wide (such as economic trends) and ϵ_{it}^m and $\epsilon_{p(i)t}^w$ represent other idiosyncratic shocks which are model-year or plant-year specific. Potential unobserved factors in δ_i include model popularity, while factors in ϵ_{it}^m could include changes in model popularity across time. Factors in $\omega_{p(i)}$ could include unobserved differences in manufacturing flexibility (including switching times) and $\epsilon_{p(i)t}^w$ may include unobserved changes in plant capabilities across time. To simplify some notation, denote $\epsilon_{it} = \epsilon_{it}^m + \epsilon_{p(i)t}^w$.

Figure 2 summarizes the theoretical factors that affect days-of-supply. Table 1 summarizes the covariates used to measure some of these theoretical factors. In the table, a negative sign (-) before the variable name indicates an inverse relationship between the theoretical factor and the variable (e.g., higher *PS* implies lower production flexibility).

We now discuss the estimation of the econometric model, (1). A primary concern is that several of the factors included in (1) may be endogenous, i.e., controlled, at least in part, by the manufacturers (e.g., the number of platforms produced at a plant). Because we do not observe all factors that affect inventory decisions, some of the endogenous variables in X and W can be correlated with the error term u. In such a situation, OLS can lead to biased estimates of β and γ in (1). For example, it is plausible that a manufacturer may assign more platforms to flexible plants (because they can better manage the additional variety). This suggests a negative correlation between $\omega_{p(i)}$ and *NPLATF.* As we have already discussed, a manufacturer may choose a higher cost markup for a popular product, which suggests that *COSTMK* and δ_i can be negatively correlated. The popularity of a model may also influence a manufacturer's decision regarding the number of options to offer, but it is not clear if this leads to more or to fewer options. The inclusion of additional controls to the model can mitigate this endogeneity bias, which we now discuss.

In our analysis, we study several specifications which include different levels of control variables. The following control variables are included in *all* of the specifications analyzed. The regressions include year indicators to control for the random component τ_t. They also include make and segment controls, which partially control for cross-sectional variation in model popularity (captured in δ_i). We used the following four-segment classification published by Wards: (i) sport cars; (ii) all other cars; (iii) Sport/Utility and Cross/Utility Vehicles (SUV); and (iv) minivans. To control for unobserved changes in model popularity across the product life-cycle (which is captured in ϵ_{it}), we include two indicators, *INTRODUM* and

ENDDUM, in the first and last year a model is produced[16]. To control for differences in replenishment leadtime, we include indicators of plant location (Mexico, Canada and U.S.) as well as a control if the model has some imported production.

Our first specification includes model indicators to control for δ_i. This is equivalent to fixed effect (FE) estimation, which exploits only the variation within each model across time. Model indicators do not entirely control for time-invariant plant unobservables, $\omega_{p(i)}$, because some models change their production across plants on different years. Hence, in the FE specifications, the estimation of γ still relies on cross-sectional variation across plants. The same applies to *PS*, which varies considerably across models produced in different plants. *SEASON*, which is time-invariant for each model, cannot be estimated with FE. We note that *DEALERS* is make-specific, hence this effect is estimated with variation across time only (in this, as well as the other specifications). The within make variation in the number of dealerships is low (the coefficient of variation is below 10% for most makes), hence we expect this effect to be estimated with low precision. *NMKT* is segment specific, so its estimate is also based on time variation only.

Our second specification is estimated without the model indicators, so it is now possible to estimate the *SEASON* effect. Here, we assume strict exogeneity, $E\left(u_{it}|X_{it}, W_{p(i)t}\right) = 0$, where X and W include all of the controls mentioned previously other than the model indicators. Given this assumption, the parameters can be estimated consistently using OLS, but random effects (RE) estimation accounts for the heteroskedastic structure of u_{it} and provides more efficient estimates. However, we note that FE estimates are consistent under less restrictive assumptions. More specifically, FE is consistent even if the assumptions $E\left(X_{it}|\delta_i\right) = 0$ and $E\left(W_{p(i)t}|\delta_i\right) = 0$ are relaxed. We use a statistical test (e.g. the Hausman test) to compare the estimates of these two specifications (and the next two) to choose a preferred one.

The third specification reintroduces the model controls and focuses on the estimation of *PS*. In particular, there may be a concern that *PS* could exhibit a mechanical relationship with the dependent variable *DS:* *PS* is evaluated with monthly inventory changes and *DS* is calculated using contemporary inventory data. To address this issue, we instrument *PS* using the following instrumental variables: the average *PS* of other models produced in the

[16]Days-of-supply in the year a model is introduced could be lower because of higher model popularity (e.g., a novel product design). Therefore, we expect *INTRODUM* to have a negative effect. We included *ENDDUM* as a control, but do not have an *a priori* predictions of the directions of its effect. We also considered further controls for product life-cycle and found no changes in our main results.

same plant (PS^{oth}), and one-year lags of the model's PS and PS^{oth}. These instruments do not use the same inventory observations, hence cannot be mechanically related to DS. They explain variation across models produced in different plants, but they are weak instruments to explain variation in PS across years within a plant. Hence, this identification strategy is not feasible when plant controls are included in the model.

Finally, our fourth specifications includes model FE and plant indicators to control for both time-invariant unobservables, δ_i and $\omega_{p(i)}$. In this specification, both β and γ are estimated using variation across years only.[17]

We report in Section 6 the results from our four main specifications. In Section 6.1 we analyze other specifications including additional controls to test the robustness of the results.

6. Results

Table 2 describes the means of the variables used, grouped by manufacturer, and some other summary statistics for the models in the sample. (We excluded some outliers from the sample, which are discussed in detail in Section 6.1.) Consistent with Figure 1, the table shows that Toyota carries approximately 30 fewer days-of-supply than the sample average. There are some other notable differences between Toyota and the other makes (primarily Chrysler, Ford and GM). Toyota has considerably higher sales per model than the other makes, substantially higher production flexibility (measured as a lower PS) and many fewer dealerships (about 1200 instead of about 3000). However, Toyota's cost markup matches the mean of the entire sample, and they are not remarkably distinctive in terms of the number of options offered per model, the number of platforms produced per plant, or plant utilization. The online appendix includes a table of correlations between the variables.

The main results are reported in Table 3. All the specifications estimate equation (1) (shown at the top of the table), but they differ in the identification strategy used, as discussed in the previous section. The bottom of the table shows whether model or plant controls (or both) are included, the number of observations, the number of models and the R-square of each regression. For the specification estimated without model controls (column (b)), the table reports the overall R-square; for the others, the within R-square is reported. To ease visualization, we do not report on the controls for year, plant location, whether the model has imports, and the *INTRO* and *END* dummy variables. Recall, the dependent

[17] For models produced in more than one plant, multiple plant indicators are set equal to one.

and independent variables are included with log transformation, so the coefficients can be interpreted as elasticities.

Column (a) reports the estimates using FE. The signs of all the point estimates are consistent with theoretical predictions (except *UTIL*, where theory is ambiguous, and *STREND+*, which is positive and not significant), but not all the coefficients are different from zero with statistical significance. To evaluate the economic significance of these results, we calculated the effect of increasing the value of the covariates one standard deviation above the mean. The number of dealerships, *DEALERS*, has the largest economic impact– an increase in this factor raises *DS* by 21%. Increasing *PS* raises *DS* by 8%. The effect of increasing *NMKT* is 9%, and raising *COSTMK* increases inventory by 6%. The effect of raising *UTIL* is 5%, and the impact of the remaining variables is below 4%.

Column (b) shows the estimates using RE, which does not control for time-invariant unobservable differences across models. Instead, it includes indicators for make and segment to partially control for these unobservables. (Make and segment are controlled via the model indicators in the other specifications.) The coefficients in columns (a) and (b) are similar with a few exceptions. The magnitude of the coefficient on *COSTMK* reduces its magnitude and becomes not statistically distinguishable from zero. The coefficient on *SALES* increases in magnitude and is negative with statistical significance. A Hausman test rejects the null hypothesis that the estimates of columns (a) and (b) are equal (p-value less than 0.01), and so the strict exogeneity assumption $E\left(X_{it}|\delta_{it}\right) = 0$ and $E\left(W_{it}|\delta_{it}\right) = 0$ is rejected by the data. A single coefficient t-test on the equality of the *COSTMK* or *SALES* coefficients also rejects the null. These results are consistent with our conjecture about the confounding effect of model popularity, as discussed previously: popular models tend to have higher sales and markups, while at the same time tend to have lower inventories for other reasons, possibly because customers are more willing to wait to purchase a popular model when it is out of stock rather than substitute to another product. Consequently, the empirical evidence suggests that controlling for model popularity is important to get consistent estimates of the direct effect of cost markup and sales volume on inventory.

Column (c) uses instrumental variables to address a possible mechanical correlation between *PS* and the dependent variable. Because the instruments include the *PS* of other models produced in the same plant and lagged values of *PS*, the sample size in this specification is smaller.[18] The standard errors increase substantially for the estimated *PS* coefficient,

[18] The sample excludes plants producing a single model and the first year in which a model is produced at

but the point estimate is similar in magnitude to (a) and significant at the 10% confidence level. The other coefficients do not change much. We estimated specification (a) over the same sample and used a Hausman test to compare the estimates. The test cannot reject that the estimates are equal. Therefore, the statistical evidence suggests that the positive effect of *PS* is not driven by a mechanical relationship with *DS*.

Column (d) includes indicators for both model and plant. Notice how the coefficient on *NPLATF* increases in magnitude and becomes significant. This provides some evidence that an increase in the number of platforms produced at a plant raises the days-of-supply of the models produced by the plant. The difference in the estimated coefficient for *NPLATF* from (a) and (d) is moderately statistically significant (p-value=.09). The other coefficients are similar in magnitude and statistical significance. This suggests that the potential bias due to unobserved plant capabilities is not large (given the controls included in our specifications).

Based on the statistical analysis, we choose (a) as our preferred specification. Specification (b) is rejected against (a), suggesting that model FE are important to control for unobservable model characteristics such as model popularity. Specification (c), which corrects for potential mechanical correlation between *PS* and *DS,* yields similar results compared to (a), but the estimates of (a) are more precise. The estimates in (d) are also similar, but model (a) is more parsimonious. In Section 6.1 we conduct additional analysis showing the robustness of the estimates of specification (a). Hence, we focus the analysis and discussion on the results provided by this specification.

The results suggests that the effect of plant utilization, *UTIL*, is positive and significant. There are multiple explanations for this effect; we defer the discussion of these explanations to Section 6.2. We also find that the number of models offered in a segment, *NMKT*, has a positive and significant effect on inventory, suggesting that increasing competition induces firms to carry higher inventories.

Our results suggest that the effect of model's sales volume (*SALES*) is not statistically significant. However, we do find that increasing the number of options in which a model is offered (*OPTIONS*) increases inventory. The magnitude of this type of demand fragmentation is small relative to the effect of increasing the number of dealerships.

We find some evidence that sales trends affect inventory levels. Negative sales trends (*STREND-*) are associated with higher inventory levels, but positive sales trends (*STREND+*)

---

a plant (which can be a new model or an exisiting model switching production between plants).

have no effect on inventory. The effect of *STREND-* could be capturing low model popularity or lags in reducing production capacity.

The effect of the number of dealerships (*DEALERS*) is found to be large, but the coefficient estimate has a large standard error. Recall that this factor is estimated using longitudinal variation in the dealership network, which is relatively low during the 9 years covered in this study. A more precise estimation of this effect would require additional time-series data or a different empirical strategy (e.g. using variation in the dealership network across geographic regions).

We used the estimates to examine how much of the difference in inventory between Toyota and Chrysler, Ford and GM is explained by *DEALERS* and *PS*. We focus in these two factors because they are economically significant and substantially different across manufacturers (see Table 2). Table 4 shows the adjusted days-of-supply for the three domestic manufacturers from setting *DEALERS* and *PS* to the average levels of Toyota, and the implied reduction in annual inventory costs. Inventory cost are calculated based on a 20% annual holding cost, $15 thousand cost per vehicle and average annual sales of each manufacturer. We also report the marginal effect of each factor and the 95% confidence interval for the adjusted days-of-supply. Recall from Table 1 that the average *DS* of Toyota is 38. The results suggest that the number of dealerships (*DEALERS*) and our measure of production flexibility (*PS*) explain almost all of the difference in days-of-supply between Toyota and Chrysler, Ford and GM.

6.1 Sensitivity Analysis

Several regression diagnostics were conducted to analyze the robustness of the results. Residuals vs. fitted scatter plots did not exhibit any systematic trend, so heteroskedasticity is not considered an issue. We found a few outliers in the data, but these are not influential points in the estimation. Excluding any observation from the data does not change any of the estimated coefficients by more than half its standard error, suggesting that the main results are not driven by influential points.

We tested alternative specifications to validate our results. (Estimation results of these alternative specifications are available from the authors upon request.) We estimated (a) without log transformations and found small differences in our results. *NMKT* and *DEALERS* are positive but not significant. The R-square is 0.3, lower than the one obtained in the log-log specification (0.38).

Four models in our sample include some imported production[19]. We excluded the model-years that included imports and found no significant change in our results. We also estimated specification (a) excluding models in their introduction year and found no changes in the main results. Recall from section 4 that our main results exclude observations from Oldsmobile in 2000-2004 and all Suzuki models due to their dramatic change in the number of dealerships. When including these observations in the analysis, the coefficient of *DEALERS* increases in magnitude and statistical significance.

Demand for more fuel-efficient vehicles increased during our sample period, possibly related to the almost 100% increase in oil prices from 1999 to 2004. To control for changes in demand across vehicle segments, we included segment-specific year controls and found no changes in our results.

We estimated specification (a) using alternative measures of *DS* as the dependent variable, based on average sales rate of one and three months ahead (instead of two months ahead). We found no change in our results, and the R-square of these specifications are also similar.

COSTMK is estimated from the data and subject to measurement error. We estimated specifications similar to (a) and (b) replacing *COSTMK* by the list price of the standard model (*PRICE*). In the FE regression, the coefficient on *PRICE* is .013 and not significant. In the RE regression, the coefficient is -0.12 and statistically significant. This change in magnitude provides further evidence of the confounding effect of model popularity. In both regressions, all the other coefficients were similar in magnitude and statistical significance. This suggests that the measurement error in *COSTMK* does not bias the estimated coefficients of the other covariates.

Our results suggest that including model FE is useful to control for unobserved model popularity to get consistent estimates of the effect of *COSTMK* on inventory performance. But if model popularity changes across time, then model FE do not control completely for this confounding effect. To test this, a proxy that captures longitudinal variation in model popularity is needed. Consumer Reports provides model ratings based on customer surveys. We included two of the measures published by Consumer Reports. The first measure is a rating from 1 to 5 based on test drives, 3 been the average rating for the segment and 5 the

[19] These include COROLLA after year 2001, and all the model-years of ACCORD, CAMRY and MAXIMA; a total of 15 observations.

highest rating[20]. The second measure is an indicator on whether the model was recommended or not. This recommendation takes into account predicted reliability (based on previous survey data) in addition to the product rating. Note that not all the models are rated each year, so the size of this sample is smaller[21]. For comparison, we estimated specification (a) using the Consumer Reports sub-sample. Adding the consumer report variables does not change the estimated coefficients. The coefficients of the consumer report variables are small and not significant. This suggests that model FE provide good controls for product popularity.

Specification (d) in Table 3 includes plant indicators to control for unobserved plant capabilities. These controls are weak if plant capabilities changed substantially over time. *PS* captures possible changes in flexibility over time, but we also included additional proxies for plant flexibility to validate our results. We obtained weekly data on work stoppages for all Chrysler, GM and Ford plants, published in Automotive News. Details of these data are described in Bresnahan and Ramey (1994).[22] From these data, we calculated the number of days that each plant was closed due to model changeovers (*MODCHG*). If a plant becomes more flexible by reducing switching times, it should be reflected in fewer plant closings (lower *MODCHG*). We estimated (a) with this additional variable. Because the sample size is smaller, *DEALERS* is no longer significant. All other estimates were similar to (a).

The specifications in Table 4 include models that were produced at more than one plant. For those models, $W_{p(i)t}$ represent average plant effects, calculated by taking the weighted mean of all plants that produced the model. To see whether this affected our results, we re-estimated specification (a) limiting the sample to models that had at least 70% of its domestic production from a single plant and included the data from that plant only in the model (the sample size reduces to 545 observations). All results were similar with two exceptions. *SALES* becomes more negative (-0.08) and statistically significant at the 10% confidence level. The coefficient on *NPLATF* is 0.08 and moderately significant (p-value<0.1). In this specification, the $W_{p(i)t}$ covariates are measured more precisely, which could explain the higher statistical significance of *NPLATF*.

[20] Consumer Report classification of vehicles includes more segments then the four we use, decomposing the car and SUV segments into multiple groups (luxury, middle/large size, etc).

[21] The sample of Consumer Reports models tends to include higher selling vehicles than in the base sample (122 versus 104 thousand vehicles).

[22] We thank Valery Ramey and Daniel Vine for providing the dataset used in their study, which includes plant closures up to 2001. We completed their dataset by collecting data from some missing plants (located in Mexico) and from years 2002-2004.

The results in Table 3 provide some evidence that the number of platforms produced at a plant affects *DS*, but the effect seems to be small. We want to test the robustness of this result with other measures of fundamental variety. A new measure was defined based on Ward's platform classification, which is different from Harbour's platform definitions.[23] We estimated specification (a) using this measure instead of *NPLATF.* The coefficient on the new measure is 0.01 with a standard error of 0.03 (and not statistically significant).

6.2 Discussion

Our two main findings are (1) fragmenting demand across more dealerships, *DEALERS*, is associated with higher days-of-supply and (2) greater production flexibility, as measured by the exhibited ability for production to track closely with sales, *PS*, is associated with lower days-of-supply. This section discusses those results as well as other issues regarding our study.

The *DEALERS* effect is of large magnitude and significant. However, it is also measured with a large standard error, which we believe is due to the limited variation in the number of dealerships across most makes over time. (Note, we control for differences across models, so *DEALERS* is not estimated with cross-sectional data.) Two makes did exhibit a considerable amount of variation, Oldsmobile and Suzuki, but we chose to exclude them from the analysis because their changes in *DS* may be due to reasons other than the shift in the number of dealerships. For example, Oldsmobile may have reduced its *DS* because it was phasing out the brand even if it was also maintaining the same number of dealerships. Furthermore, although we are able to identify an important effect regarding the dealership network, we are unable able to identify the precise mechanism by which the number of dealerships is related to *DS*. For example, it is possible that increasing the number of dealerships leads each dealership to have more variable demand, thereby requiring each dealer to carry a higher *DS* to maintain the same service level (such as a fill rate target or in-stock probability). Alternatively, increasing the number of dealerships may increase the amount of competition the brand faces because the new dealerships will be closer to existing dealerships from the same brand as well as dealerships from other brands. Theory is ambiguous on the impact of competition on inventory, but some models suggest that increased competition can increase

[23] Ward's assigns more than one platform to some models during a calendar year. For example, they considered several platforms for the Toyota Camry, so that the first half of a calendar year the Camry was produced in one platform and on the second half, after the model change-over, on another platform. This suggests that their platform classification is more sensitive to minor changes in the model specifications.

a firm's optimal inventory. Therefore, we cannot distinguish with our data if the *DEALERS* effect is due to demand fragmentation or increased competition or a combination of both. Interestingly, our other measures of demand fragmentation do not suggest a strong effect. For example, we did not find a significant economies of scale (DS is not associated with higher or lower sales per model), and the effect of the number of options offered for the model is small (but still statistically significant). It is possible that economies of scale are adequately captured by our other controls. For example, if PS is removed from the regression, the effect of *SALES* increases in magnitude and becomes significant. The option effect may be small due to conflicting forces: adding options may fragment demand and make demand more variable, but product differentiation offers better match to heterogeneous customer preferences, making each option more popular and thereby require less inventory. (See Cachon et al. (2006) for a model of some of these effects.) Furthermore, there is evidence in the literature that the number of options may not have a strong effect on production (Fisher and Ittner (1999)).

We also find an important association between our proxy of production flexibility, PS, and our dependent variable, DS. PS measures inventory volatility and we suggest that more flexible plants generate less inventory volatility because they are able to better match their production to their demand. Furthermore, we suggest that inventory volatility can vary independently of the average inventory level but firms with lower inventory volatility tend to carry less inventory, possibly because their lower inventory volatility enables them to choose to operate with lower safety stocks. Consistent with a connection between PS and production flexibility, we find a higher correlation between the PS of models produced at the same plant than between models in the same segment. However, we acknowledge that a concern can remain that there exists a mechanical relationship between PS and DS. Hence, we estimated a specification using instrumental variables for which there cannot be a mechanical relationship between DS and PS (e.g., we use the PS of other models produced at the same plant to instrument for a model's PS). We continue to observe this relationship with this specification, which provides additional support for our hypotheses that PS proxies for production flexibility. Furthermore, PS appears to be capturing a measure of production flexibility beyond just the number of platforms produced at a plant, *NPLATF,* or the aggregate scale of production, *SALES*. However, with our data we are unable to identify the specific mechanism that enables one model's production to track sales more closely than another model's production. For example, PS could reflect lower switching times or more flexible labor, among other possible sources of production flexibility.

We find that higher cost markups are associated with higher inventory, which provides evidence of the direct effect of markups on shortage costs. However, our econometric analysis also suggests that unobservable model characteristics (such as model popularity) can confound the direct effect of markups on inventory. Popular models tend to have higher markups, but may also have lower inventories because customers are more willing to wait for the product. Hence, a regression that does not include controls for model popularity may underestimate the direct positive effect of cost markups on inventory. This appears to be an important issue for the auto industry and may be relevant for other industries as well.

Our results suggest that models produced at highly utilized plants have higher inventories. Two alternative explanations are consistent with this finding. The first one is that highly utilized plants may have longer production leadtimes, which leads to higher safety stocks. The second explanation is related to fixed plant production capacity. In plants producing more than one product with a cyclic schedule, switching times reduce effective capacity available for production. To meet an increase in demand with fixed capacity, plants need to schedule longer production cycles, which increases production lot sizes and utilization (because production volume increases and capacity is fixed). Consequently, higher plant utilization is associated with higher inventory levels. Since we do not have data on production leadtimes and lot sizes, we cannot identify these two effects separately.

We also find that competition, measured by the number of models offered in the same segment, has a positive effect on inventory. When more substitutes are available, customers may be less prone to wait for a product that is out of stock. Consequently, stronger competition could make stock-outs more costly to a firm, leading to higher target service levels (and thereby higher inventories) to reduce the frequency of stock-outs.

7. Conclusion

We report substantial and persistent differences in finished-goods inventory levels in the U.S. auto industry: data on days-of-supply suggest that Toyota's well documented advantage in manufacturing efficiency and upstream supply chain management extends to their finished goods supply chain downstream from their assembly plants to their dealerships. We identify and measure the effect of several factors on finished-goods inventory in this industry. We find that two of these factors, production flexibility and the number of dealerships, explains most of the difference in inventory between Toyota and Chrysler, Ford and GM. (Although we use Toyota for our benchmark for comparisons, our qualitative results are similar for Honda.) Production flexibility allows a firm to track production more closely to sales, thereby yielding a lower optimal level of safety stock for a firm. Fewer dealerships allows a firm to pool demand in fewer locations and to reduce both intra brand and inter-brand competition, either of which or both could lead to a lower optimal inventory level. Furthermore, we find the dealership effect to be the most influential: e.g., this factor alone explains more than 75% of the difference in inventory between Toyota and GM.

While it is debatable whether other manufacturers can emulate Toyota's skill at production flexibility, it is clear that it will be difficult for firms like the established domestic producers to match Toyota's advantage in terms of its dealership network. Chrysler, Ford and GM established their dealership networks in the first half of the 20^{th} century, before the inter-state highway system and at a time when the U.S. was more rural. As a result, they created many dealerships so that consumers need not travel far to reach a dealer. Toyota (and other later entrants to the U.S. market, like Honda) did not need to open nearly as many dealerships because as transportation became easier, consumers were willing to travel farther (or did not need to travel as far with increased urbanization). Furthermore, because the franchise laws in most states impose stringent requirements on the opening and closing of dealerships, firms like GM are unable to easily change their dealership network, either the number of dealerships or their locations. For example, during the phase-out of the Oldsmobile brand during 2001-2004, GM spent more than $1 billion reimbursing dealers for forgone profits and equipment (Welch (2006)); and Ford attempted to consolidate dealerships in local markets, but they found the legal barriers to be insurmountable (Warner (1998)). Thus, it appears that Toyota has a competitive advantage in the U.S. with respect to finished-goods inventory that cannot easily be eliminated.

Although our results may appear to argue for more flexibility in franchise laws, we caution against such a quick conclusion. Changes in those laws should be based, at least in part, on how they would effect consumer welfare and our results suggest conflicting effects. Consumer welfare should increase if there is more inventory available to choose from, but it is not clear how increasing the number of dealerships will influence availability: holding the sale rate constant, our results suggest that days-of-supply will increase, which should increases the absolute number of vehicles at dealerships; but an increase in dealerships could lower the sales per dealership, which, holding the days-of-supply constant, could lead to fewer vehicles available to consumers. Furthermore, more dealerships should increase price competition, which is beneficial to consumers, but more dealerships also raises inventory carrying costs, which harms consumer welfare. Therefore, a detailed analysis of the magnitude of these effects is needed before a clear recommendation can be made regarding changes to franchise laws.

References

Berry, Steve, James Levinsohn, Ariel Pakes. 1995. Automobile prices in market equilibrium. *Econometrica* **963(4)** 841–890.

Bresnahan, T.F., V.A. Ramey. 1994. Output fluctuations at the plant level. *Quarterly Journal of Economics* **109**(3) 593–624.

Cachon, Gerard P., Martin Lariviere. 1999. Capacity Allocation Using Past Sales: When to Turn-and-Earn. *Management Science* **45**(5) 685–703.

Cachon, Gerard P., Taylor Randall, Glenn Schmidt. 2007. In Search of the Bullwhip Effect. *Manufacturing & Service Operations Management* **9**(4) 457.

Cachon, Gerard P., Christian Terwiesch, Yi Xu. 2006. The impact of consumer search, firm entry and competition on assortment and pricing. *Marketing Science* Forthcoming.

Clark, Kim, Takahiro Fujimoto. 1989. Lead time in automobile product development: Explaining the japanese advantage. *Journal of Engineering and Technology Management* **6** 25–58.

Fisher, Marshall L., Christopher Ittner. 1999. The impact of product variety on automobile assembly operations: Empirical evidence and simulation analysis. *Management Science* **45**(6) 771–786.

Gaur, Vishal, Marshall L. Fisher, Ananth Raman. 2005. An econometric analysis of inventory turnover performance in retail services. *Management Science* **51**(2) 181–194.

Goyal, Manu, Serguei Netessine, Taylor Randall. 2006. Deployment of manufacturing flexibility: an empirical analysis of the North American automotive industry. Working paper, http://www.netessine.com/.

Hall, George. 2000. Non-convex costs and capital utilization: A study of production scheduling at automobile assembly plants. *Journal of Monetary Economics* **45** 681–716.

Harris, Donna. 2004. Public retail groups chop inventories. Automotive News, November 22, 2004.

Hendricks, Kevin B., Vinod R. Singhal. 2005. Association between supply chain glitches and operating performance. *Management Science* **51**(5) 695–711.

Kekre, Sunder, Kannan Srinivasan. 1990. Broader product line: A necessity to achieve success? *Management Science* **36**(10) 1216–1231.

Lai, Richard. 2006. Inventory signals. Harvard TOM Research Paper Series No. 05-15, http://ssrn.com/abstract=723381.

Lieberman, Marvin, Shigeru Asaba. 1997. Inventory reduction and productivity growth: A comparison of japanese and us automotive sectors. *Managerial and Decision Economics* **18**(2) 73–85.

Lieberman, Marvin, Lawrence Lau, Mark Williams. 1990. Firm-level productivity and management influence: A comparison of u.s and japanese automobile producers. *Management Science* **36**(10) 1193–1215.

Lieberman, Marvin B., Lieven Demeester. 1999. Inventory reduction and productivity growth: Linkages in the Japanese automotive industry. *Management Science* **45**(4) 466–485.

MacDuffie, John Paul, Kannan Sethuraman, Marshall Fisher. 1996. Product variety and manufacturing performance: Evidence from the international automotive assembly plant study. *Management Science* **42**(3) 350–369.

Marti, Eric. 2000. Desintermediation in the U.S. auto industry. Stanford case EC-10.

Nahmias, Steven. 2005. *Production and Operations Management*. 5th ed. McGraw-Hill Irwin, New York, NY.

Nevo, Aviv. 2000. A practioner's guide to estimation of random-coefficiets logit models of demand. *Journal of Economics and Management Strategy* **9**(4) 531–548.

Olivares, Marcelo, Gerard P. Cachon. 2007. Competing retailers and inventory: An empirical investigation of U.S. automobile dealerships. SSRN working paper, http://ssrn.com/paper=993796.

Rajagopalan, S., Arvind Malhotra. 2001. Have U.S. manufacturing inventories really decreased? an empirical study. *Manufacturing & Service Operations Management* **3** 14–24.

Randall, Taylor, Serguei Netessine, Nils Rudi. 2006. An empirical examination of the decision to invest in fulfillment capabilities: A study of internet retailers. *Management Science* **52**(4) 567–580.

Rumyantsev, Serguei, Serguei Netessine. 2007. What Can Be Learned from Classical Inventory Models? A Cross-Industry Exploratory Investigation. *Manufacturing & Service Operations Management* **9**(4) 409.

van Ryzin, Garrett, Siddarth Mahajan. 1999. On the relationship between inventory costs and variety benefits in retail assortments. *Management Science* **45**(11) 1496–1509.

Warner, Fara. 1998. Ford motor reaches agreement to create retail network in Utah. Wall Street Journal, May 19.

Welch, David. 2006. The other club battering GM. *Businessweek* **3974**(38).

Wu, Owen Q., Hong Chen, Murray Z. Frank. 2005. What actually happened to the inventories of American companies between 1981 and 2000. *Management Science* **51**(7) 1015–1031.

Data Sources

The Harbour Report, Harbour Consulting, 1996-2004.

Ward's AutoWorld, Ward's Auto, 1994-2004

Automotive News, Crain Communications

Consumer Reports, Consumer Union, 1996-2004

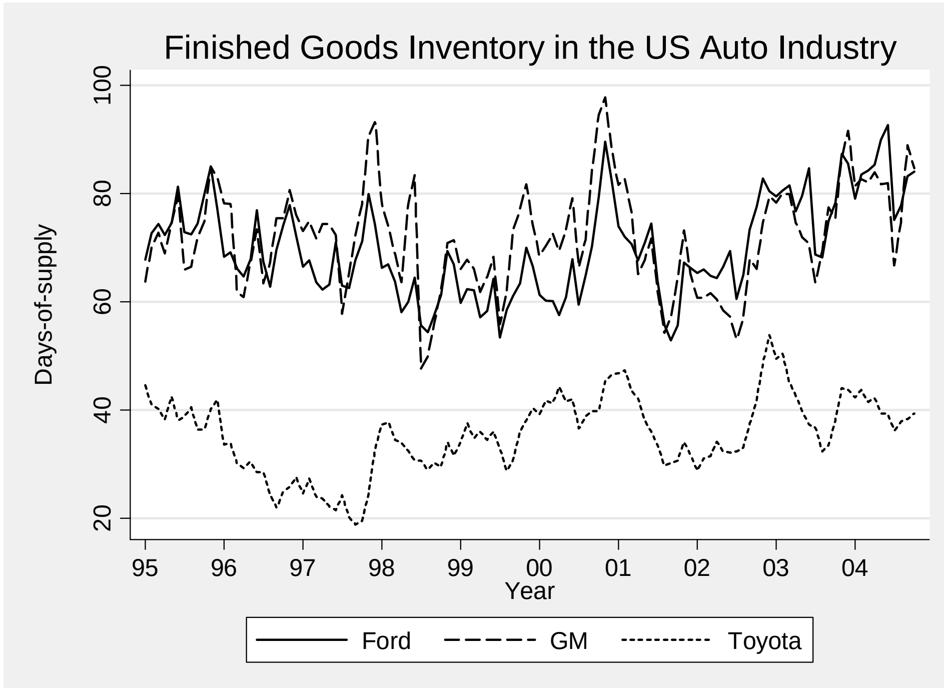

Figure 1 – Inventory of three auto manufacturers. Days-of-supply is calculated as the aggregate inventory at the end of each month divided by the average daily sales rate in the following two months. Inventory includes all finished vehicles in US territory, including inventory in the plant, in ports of entry, in transit to dealers and in dealerships.

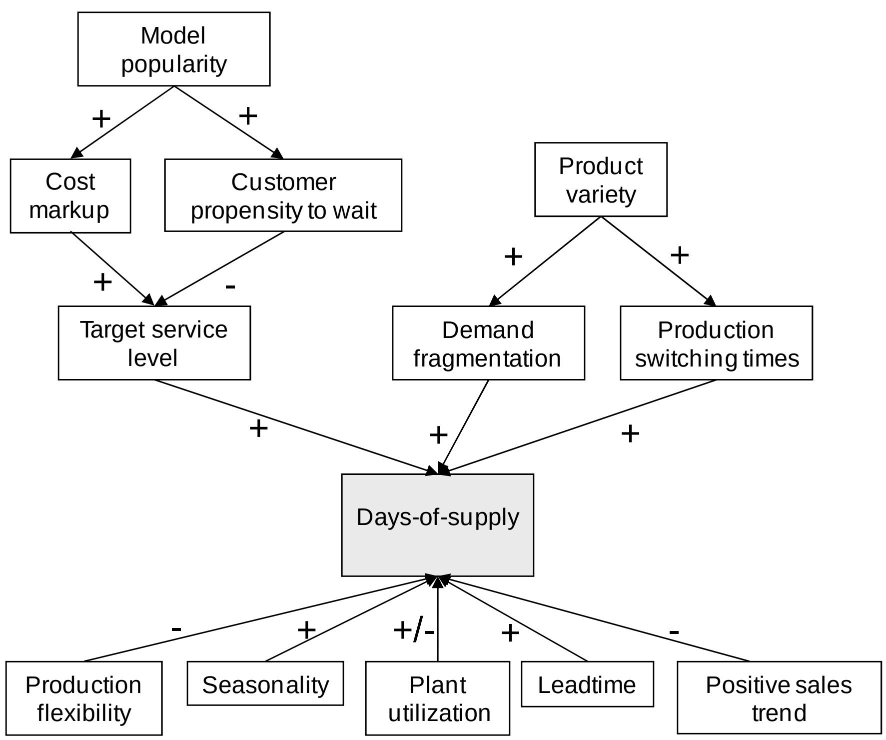

Figure 2 – Theoretical factors affecting days of supply.

Covariate	Definition	Factors
NMKT	# of models in the same segment	Competition
COSTMK	Estimated cost markup	Cost markup, model popularity
SALES	Average monthly sales	Demand fragmentation (-), model popularity
OPTIONS	# of model options	Demand fragmentation
DEALERS	# of dealerships	Demand fragmentation, customer propensity to wait (-)
NPLATF	# of platforms produced in plant	Production switching times, production flexibility
PS	Average difference of production and sales (equation (4))	Production flexibility (-)
SEASON	Seasonal variation (equation (3))	Seasonality, model popularity (-)
UTIL	Total plant production divided by capacity	Plant utilization
STREND+	Positive sales trend (equation (2a))	Positive sales trend, model popularity
STREND-	Negative sales trend (equation (2b))	Positive sales trend (-), model popularity (-)

Table 1 – Summary of the covariates included in the econometric analysis. A negative sign (-) after the factor name indicates an inverse relationship between the factor and the variable. For example, higher *PS* implies lower production flexibility.

Manufact.	DS	SALES	STREND-	STREND+	DEALERS	OPTIONS	COSTMK	NMKT	UTIL	NPLATF	PS	SEASON
Chrysler	69	9147	620	660	2884	6.16	0.68	76	0.40	1.18	0.21	0.16
Ford	74	10499	411	839	3056	6.44	0.69	89	0.38	1.53	0.22	0.14
GM	77	7611	648	562	3075	4.80	0.63	85	0.36	1.37	0.25	0.15
Honda	45	12949	762	206	582	4.70	0.60	104	0.41	1.30	0.19	0.12
Nissan	80	7725	448	417	1076	4.15	0.62	78	0.36	1.54	0.25	0.15
Toyota	38	16473	879	1169	1197	4.87	0.65	100	0.42	1.27	0.13	0.12
mean	72	9066	587	653	2685	5.43	0.65	86	0.37	1.39	0.23	0.15
sd	24	7982	1388	1202	1177	4.54	0.10	41	0.11	0.67	0.11	0.05
min	13	281	0	0	258	1.00	0.23	11	0.06	1.00	0.06	0.07
max	197	37347	14540	16376	4420	38.00	1.02	124	0.64	4.00	1.03	0.32

Table 2 – Summary statistics (variables measured without log transformation). The means of the variables are also reported separately for the six major manufacturers.

Model: $$DS_{it} = \beta X_{it} + \gamma W_{p(i)t} + u_{it}$$

	(a)	(b)	(c)	(d)
SALES	-0.016	-0.105**	-0.001	-0.054
	(0.040)	(0.024)	(0.047)	(0.045)
STREND-	0.019**	0.016*	-0.002	0.019**
	(0.006)	(0.006)	(0.007)	(0.006)
STREND+	0.005	0.002	-0.015*	0.006
	(0.006)	(0.006)	(0.007)	(0.006)
DEALERS	0.528*	0.483*	0.699**	0.452*
	(0.207)	(0.210)	(0.237)	(0.218)
OPTIONS	0.058*	0.072**	0.016	0.057*
	(0.024)	(0.021)	(0.026)	(0.025)
COSTMK	0.439**	0.191^	0.311^	0.554**
	(0.154)	(0.110)	(0.171)	(0.162)
NMKT	0.220**	0.213**	0.266**	0.189*
	(0.071)	(0.065)	(0.086)	(0.079)
UTIL	0.178**	0.129**	0.211**	0.226**
	(0.041)	(0.035)	(0.046)	(0.045)
NPLATF	0.047	0.073*	0.047	0.088^
	(0.042)	(0.033)	(0.054)	(0.052)
PS	0.194**	0.218**	0.213^	0.175**
	(0.027)	(0.026)	(0.111)	(0.027)
SEASON		0.216**		
		(0.056)		
Model controls	Yes	No	Yes	Yes
Plant controls	No	No	No	Yes
# obs	705	705	600	705
# models	133	133	122	133
R-squared	0.39	0.60	0.37	0.44

Table 3 – Main estimation results. Standard errors showed between parentheses. All the covariates are included with log transformation. Column (c) uses instrumental variables to instrument for PS (using PS of other models produced at the same plant and lagged values of PS as instrumental variables). Column (b) is estimated with RE; all other specifications are estimated with model FE. Column (d) also includes plant indicators. ^,* and ** indicate statistical significance at the 0.1, 0.05 and 0.01 confidence levels, respectively.

Manuf.	Days of Supply	% reduction in days-of-supply PS	DEALERS	Adjusted Days of Supply Estimate	95% CI	Inv. Cost Reduction M$
Chrysler	69	7.9%	35.6%	41	[27 , 55]	$402
Ford	74	9.2%	37.4%	42	[27 , 58]	$638
GM	78	11.6%	37.3%	43	[28 , 60]	$957

Table 4 - Reduction in days-of-supply and inventory costs (in million $ per year) for Chrysler, Ford and GM from adjusting production flexibility (PS) and the number of dealerships (DEALERS) to the average levels of Toyota. For adjusted days-of-supply, the point estimate and 95% confidence interval (CI) are reported. Inventory costs were calculated based on $15,000 cost per vehicle and 20% annual holding cost.

ONLINE APPENDIX

Estimation of the Cost Markups

We used the methodology described in Berry et. al.(1995) (hereon BLP) to estimate markups for each model on each year. We briefly outline this methodology and refer interested readers to the original articles for details.

This methodology forms part of a broad class of econometric models known as structural estimation. Assuming a theoretical model of decision making, this approach attempts to impute the parameters of the decision model that best fit the observed data. In our application, following BLP, we assume the U.S. automotive industry is an oligopoly where auto prices are determined through Bertrand/Nash equilibrium in a differentiated product market. If we knew the complete specification of demand for automobiles and the marginal costs of manufacturing each car, then we could compute the equilibrium sales and prices (and therefore the markups) for each model. Structural estimation follows the reverse approach. Given the equilibrium prices and market shares observed in the data, we estimate the costs and demand parameters that are consistent with this equilibrium under the maintained behavioral assumption of the agents in the market.

Let p_j and c_j be the price and (constant) marginal cost of model j in a specific market. Let J_f be the set of models offered by firm f. Demand is defined by the function $D_j\left(\vec{p},\vec{x}\right)$ which specifies the total demand of product j as a function of the vector of prices $\vec{p}$ and the vector of characteristics $\vec{x}$ of all the products offered in the market. The profit function for firm f is given by:

$$\pi_f = \sum_{j\in J_f} \left(p_j - c_j\right) D_j\left(\vec{p},\vec{x}\right) - \text{Fixed Costs}_j$$

Denote by M the total market size and $s_j\left(\vec{p},\vec{x}\right) = D_j\left(\vec{p},\vec{x}\right)/M$ the market share of model j. Under Nash equilibrium in prices, the equilibrium price vector $\vec{p}^*$ and market share vector $\vec{s}^*$ must satisfy:

$$\vec{s}^* + \sum_{r\in J_f} \left(p_r^* - c_r\right) \times \frac{\partial s_r\left(\vec{p}^*,\vec{x}\right)}{\partial p_j} = 0 \tag{1}$$

for all models j in the market. Denote the equilibrium markup of model j by $b_j^* = p_j^* - c_j$. Let b^* be the stacked vector of these markups and define the matrix Δ of price substitutions by $\Delta_{jr} = \partial s_r/\partial pr$ if product j and r are produced by the same firm, and $\Delta_{jr} = 0$ otherwise.

Equation (1) can be rewritten in matrix form as:

$$b^* = \Delta^{-1} s^* \tag{2}$$

where b and s are the stacked vectors of markups and market shares, respectively.

Equation (2) specifies the markups as a function of the price substitutions and the equilibrium market shares. Therefore, if we knew Δ, we could compute the markups using the observed market shares, assuming these quantities were reached through the Nash equilibrium in prices described above. The matrix Δ is defined by: (i) the set of models produced by each firm, which we observe in the data; and (ii) the price substitutions $\partial s_j / \partial p_r$, which we do not observe directly. Therefore, in order to compute the markups, we need to estimate demand in the auto industry. As summarized by Nevo (2000), researchers face two main difficulties in estimating demand for differentiated products: (i) the number of substitutions to be estimated increases in the square of the number of products; and (ii) prices are endogenous and tend to be correlated with unobserved product quality.

To overcome the first problem, BLP uses a random-coefficient multinomial logit which incorporates unobserved heterogeneity in consumer tastes for product attributes, allowing for flexible substitution patterns in the data. To overcome the second problem, BLP uses instrumental variables. Valid instruments for price are cost shifters and the characteristics of *other products* available in the market.

We estimated demand using data from 1996 to 2004. We defined the U.S. market sales for a given year as our market, giving a total of 9 markets. As in BLP, we used the number of households in the U.S. as our size of the market. We included all car models, all SUVs and CUVs and minivans. We excluded pickups and full size vans from the analysis. In addition to prices, we included random coefficients for the following product characteristics: (i) car size (the product of the length and width of the vehicle); (ii) a measure of acceleration (horsepower divided by weight); (iii) Miles per dollar of gasoline (based on EPA Miles per gallon measures and the average price of gasoline during the year); (iv) a measure of security (the sum of the indicators on whether the base model has airbags and Automatic Break System); and (v) indicators on market segment (whether the auto was classified in the following non-overlapping segments: cars, sport cars, SUV and minivan). We let the price coefficient interact with personal income, which was drawn from a lognormal distribution. The median income varies across years and was obtained from CPS. All dollar values are normalized to

1982-83 dollars using the CPI. We included non-random coefficients in the demand side to control for brand (Ford, GM, Chrysler, Honda, Toyota, Nissan, other European brands and other Asian brands). We added two measures based on the ratings from Consumer Reports: a dummy to indicate whether the model was recommended and a dummy on whether the product was rated above average. On the supply side, we included all the product characteristics (without the brand dummies nor the Consumer Report variables) plus a trend. We added variables indicating the percentage of production coming from plants in the U.S., Canada and Mexico and from overseas[1]. We did not include wages and exchange rates in the cost side as in BLP.

Our implementation was based on the source code provided by Nevo (2000). We modified the source code to include the supply equation into the estimation. We also added routines to compute the optimal instruments described in BLP. Both the supply side and the use of optimal instruments improved significantly the precision of our coefficient estimates. Overall, our results are consistent with those obtained in BLP, especially in the estimation of the price elasticities and cross-substitutions.

Table A3 provides a sample of the estimated markups for different years. The markups look reasonable both in absolute magnitude and in relative comparison across models. "Mainstream" cars, such as the Honda Accord, the Toyota Camry and the Ford Taurus tend to have smaller markups, while Sport models (Porsche 911), SUV's (Chevy Suburban) and expensive luxury vehicles (Mercedes E Class and Jaguar XJ6) tend to have larger markups. Even though the markups are increasing with the price of a vehicle, markups and cost markups are not proportional. Note also that some models exhibit unreasonable markups. For example, the Mini Cooper exhibits very low markups because we do not capture specific attributes of these exclusive models in our product characteristics. The sports car segment also exhibits particularly low cost markup. We also observe that there is some variation in the markups across time within a model, which is useful to identify the effect of this variable on inventory performance. Table A4 shows the price semi-elasticities, defined as the percent change in market share after a $1000 price change, for a sample of models. The table shows that

[1] We further classified imports into Europe, Japan and other Asian countries based on their make. For models that have both domestic and imported production, we aproximated the percentage of imports by the difference in annual sales and domestic production.

price changes tend to have the largest effect on products in the same segment[2]. We also note that the magnitude of the semi-elasticities differ across segments, which suggests that the propensity to substitute differs across vehicle types. This provides further support for controlling for market segment in our econometric model.

To validate our cost markup estimates, we calculated the gross profit per vehicle for each make implied by the model markups and compared it to the dealership gross profit per vehicle published by Automotive News. If dealerships get a fixed proportion of the supply chain profits and we measure the markups precisely, this two measures should be perfectly correlated. The correlation between the two measures is 0.8 (approximately). A regression with dealership profit as the dependent variable and the calculated markups and an intercept as covariates gives a coefficient of determination (R^2) of 0.7, that is, 70% of the variation in dealership gross profits can be explained by our estimated markups.

[2] If the choice model were a standard Multi-nomial logit (MNL) model, all the rows in this table would then be identical. Therefore, including segment indicators as random coefficients helps to overcome the restrictions of the MNL.

ADDITIONAL TABLES

Manufacturer	Total Domestic	No. Plants Excluded	Excluded Plants	Products in excluded plants
AM General	2	2	All	
AutoAlliance	1	0		
BMW	2	2	All	
CAMI	1	0		
Chrysler Corp.	16	1	Conner	Dodge Viper, Prowler
Ford N.A. Mfg.	23	0		
General Motors	34	0		
Honda	6	1	Lincoln	Odyssey(*), Pilot(*)
Mercedes Benz	2	2	All	
Mitsubishi	1	0		
Nissan	4	1	Canton	Infiniti QX56, Altima(*), Armada, Quest, Titan
NUMMI	2	0		
Subaru	1	1		
Toyota	4	2	Princeton, Tijuana	Sequoia, Sienna(*), Tundra, Tacoma
Volkswagen	1	1	All	
Volvo	1	1	All	

(*) Models were also produced at other plants included in our dataset.

Table A1 - Detail of domestic plants of each manufacturer which *do not* have data available in the Harbour Report.

Manufacturer	1996	1997	1998	1999	2000	2001	2002	2003	2004	Total
AM General					0%	0%	0%	0%	0%	0%
AutoAlliance	0%	100%	100%	100%	100%	100%	100%	100%	100%	87%
BMW	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
CAMI	58%	100%	100%	100%	100%	100%	100%	100%	100%	94%
Chrysler Corp.	100%	100%	100%	100%	100%	100%	100%	100%	94%	99%
Ford N.A. Mfg.	100%	98%	100%	100%	100%	100%	100%	100%	100%	100%
General Motors	97%	98%	98%	99%	100%	96%	96%	95%	95%	97%
Honda	81%	79%	79%	71%	66%	64%	73%	70%	67%	72%
Mercedes Benz	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Mitsubishi	0%	0%	0%	100%	100%	100%	100%	100%	100%	66%
Nissan	75%	70%	62%	64%	55%	50%	55%	55%	45%	57%
NUMMI	100%	100%	100%	100%	100%	100%	98%	100%	100%	100%
Subaru	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Toyota	100%	100%	100%	92%	84%	78%	79%	68%	61%	82%
Volkswagen	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Volvo	0%	0%	0%							0%
TOTAL	91%	92%	92%	92%	91%	88%	89%	87%	85%	90%

Table A2 – Percentage of total domestic production covered by our sample, by manufacturer by year.

Model	Avg Markup	Avg CostMk	2001	2002	2003	2004
ACCENT	3.964	62%	65%	66%	71%	55%
NEON (DODGE)	4.666	54%	56%	57%	58%	56%
SENTRA	4.930	58%	57%	58%	60%	58%
CIVIC	5.079	63%	63%	64%	67%	65%
FOCUS	5.433	68%	75%	63%	66%	63%
ESCORT	5.503	68%	61%	74%		
MINI COOPER	6.118	54%		55%	53%	55%
GOLF	6.219	61%	62%	60%	62%	61%
RAV4	7.174	65%	63%	61%	63%	55%
CAMARO	7.304	61%	58%	59%		
STRATUS SEDAN	7.464	67%	69%	68%	69%	66%
ACCORD	7.470	65%	67%	67%	62%	61%
MALIBU	7.477	66%	62%	61%	66%	75%
CAMRY	8.115	69%	74%	63%	66%	64%
SEBRING COUPE	8.277	65%	68%	69%	68%	66%
IMPALA	8.720	68%	67%	66%	69%	70%
MONTE CARLO	8.731	68%	66%	65%	66%	68%
CORVETTE	8.970	26%	27%	25%	24%	25%
TAURUS	9.097	77%	77%	76%	77%	72%
CARAVAN	9.928	92%	98%	99%	103%	85%
SAAB 9-3	10.278	56%	56%	54%	60%	56%
ENVOY	10.539	54%		53%	53%	56%
VOLVO 60	10.773	61%	59%	59%	60%	65%
EXPLORER	10.793	75%	73%	72%	75%	68%
WINDSTAR PASS	10.959	83%	84%	81%	82%	
CHRYSLER 300M	11.471	57%	58%	55%	56%	
GRAND CHEROKEE	11.561	65%	66%	64%	64%	61%
TOWN & COUNTRY	12.169	75%	84%	77%	79%	81%
CHEVY SUBURBAN	12.863	69%	64%	63%	64%	68%
VIPER	12.981	20%	25%	25%	24%	25%
AUDI A6	14.434	61%	61%	58%	59%	61%
MERCEDES M CLASS	15.098	63%	64%	64%	64%	60%
BMW X5	15.913	63%	64%	63%	61%	62%
BMW 5 SERIES	16.657	65%	64%	63%	65%	69%
LEXUS GS300	17.201	63%	63%	62%	62%	62%
ACURA RL	18.546	65%	65%	63%	64%	63%
MERCEDES E CLASS	20.389	70%	67%	67%	67%	65%
JAGUAR XJ6/8	23.897	62%	62%	61%	64%	59%
PORSCHE 911	27.286	59%	62%	57%	61%	62%

Table A3 – Estimated markups (price – cost) and cost markups (markup/cost). Markups are in thousand dollars of 2004.

	ACCENT	CIVIC	CAVALIER	ACCORD	CAMRY	MALIBU	TAURUS	SIENNA	CARAVAN	EXPLORER	GRAND CHEROKEE	VIPER	MERCEDES E CLASS	PORSCHE 911
ACCENT	-27.49%	0.44%	0.29%	0.46%	0.62%	0.19%	0.16%	0.07%	0.07%	0.13%	0.06%	0.00%	0.01%	0.00%
CIVIC	0.06%	-19.79%	0.35%	0.36%	0.27%	0.28%	0.25%	0.08%	0.14%	0.12%	0.07%	0.00%	0.01%	0.00%
CAVALIER	0.07%	0.56%	-19.72%	0.39%	0.27%	0.31%	0.28%	0.08%	0.14%	0.13%	0.07%	0.00%	0.01%	0.00%
ACCORD	0.05%	0.29%	0.20%	-13.87%	0.50%	0.21%	0.19%	0.10%	0.10%	0.17%	0.09%	0.00%	0.03%	0.00%
CAMRY	0.06%	0.20%	0.12%	0.45%	-13.83%	0.14%	0.13%	0.08%	0.05%	0.18%	0.09%	0.00%	0.04%	0.00%
MALIBU	0.03%	0.33%	0.23%	0.30%	0.22%	-13.28%	0.23%	0.08%	0.14%	0.13%	0.07%	0.00%	0.01%	0.00%
TAURUS	0.03%	0.31%	0.22%	0.29%	0.22%	0.25%	-12.48%	0.08%	0.14%	0.13%	0.07%	0.00%	0.01%	0.00%
SIENNA	0.02%	0.15%	0.10%	0.23%	0.23%	0.13%	0.12%	-10.22%	0.64%	0.13%	0.07%	0.00%	0.02%	0.00%
CARAVAN	0.01%	0.18%	0.12%	0.15%	0.10%	0.15%	0.14%	0.42%	-9.77%	0.10%	0.05%	0.00%	0.01%	0.00%
EXPLORER	0.02%	0.11%	0.07%	0.19%	0.23%	0.10%	0.10%	0.06%	0.07%	-9.69%	0.18%	0.00%	0.02%	0.00%
GRAND CHEROKEE	0.02%	0.11%	0.08%	0.19%	0.22%	0.10%	0.10%	0.06%	0.07%	0.34%	-9.60%	0.00%	0.02%	0.00%
VIPER	0.00%	0.07%	0.08%	0.15%	0.08%	0.06%	0.06%	0.05%	0.03%	0.07%	0.05%	-6.15%	0.07%	0.07%
MERCEDES E CLASS	0.01%	0.05%	0.03%	0.18%	0.29%	0.06%	0.06%	0.05%	0.04%	0.10%	0.06%	0.00%	-5.40%	0.01%
PORSCHE 911	0.00%	0.02%	0.02%	0.09%	0.10%	0.03%	0.03%	0.05%	0.03%	0.06%	0.04%	0.01%	0.08%	-3.80%

Table A4 – Estimated price semi-elasticities. Each entry indicates the percent change in sales of the row model after a $1000 change in price of the column model (price changes in 2004 dollars).

	DS	SALES	UTIL	PS	DEALERS	NMKT	COSTMK	NPLATF	STREND+
DS	1.000								
SALES	-0.133	1.000							
UTIL	-0.123	0.293	1.000						
PS	0.273	-0.232	-0.290	1.000					
DEALERS	-0.115	0.198	0.188	-0.151	1.000				
NMKT	0.223	-0.023	-0.280	0.082	-0.239	1.000			
COSTMK	-0.034	0.115	0.024	-0.079	0.274	-0.304	1.000		
NPLATF	0.152	-0.129	-0.033	0.099	-0.140	0.034	-0.106	1.000	
STREND+	-0.002	0.003	0.000	-0.013	0.001	-0.001	0.001	0.002	1.000
STREND-	-0.008	-0.005	0.000	0.010	0.005	-0.003	0.003	0.002	-0.056

Table A5 – Correlation matrix. Shows the correlations *within* each panel, after controlling for model FE.