

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

October 24, 2006; Revised January 31, 2008, October 7, 2008, June 3, 2009

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.

*For their helpful comments, the authors thank the Associate Editor, three reviewers and seminar participants at Northwestern University, University of California at Los Angeles and University of California at San Diego. 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 produce design and manufacturing technology (e.g., the assembly line, just-in-time inventory, kan-ban, etc.). 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 brands – Chevrolet, Ford and Toyota– between 1995 and 2004. This measure of inventory performance includes all non-truck 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 Chevrolet 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 persistent¹. 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. For example, theory predicts that fragmenting demand across different products (i.e., vehicle models), or across different options of a given product or across different geographic locations (e.g., dealerships) leads to more variable demand and therefore more inventory. Heightened competition can influence a firm’s inventory in at least two ways: (1) it should reduce a product’s margin, which leads to lower inventory; and/or (2) it gives consumers

¹We 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).

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. Production capabilities should also influence inventory. As a plant becomes more flexible over time it can adjust its production more readily and therefore can better match its production to its sales. Hence, adding flexibility to a plant could enable a firm to hold less inventory. 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. 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).

Among our findings, we highlight our observation that a manufacturer's inventory is associated with (i) the number of dealerships in a manufacturer's distribution network and (ii) the level of production flexibility the manufacturer exhibits. In particular, inventory reductions are related to decreases in dealerships and increases in production flexibility. In fact, these two factors appear sufficient to explain the differences in finished goods inventory between Toyota and three other makes, Chrysler, Ford and General Motors.

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 our hypotheses. Section 6 details our estimation methods, provides our estimation results and offer sensitivity analysis. Section 7 summarizes our main conclusions.

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 obtain 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 during the time of 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 pur-

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

chased directly from dealership inventory³. 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)). 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 are 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,

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

⁴Exports 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.

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

Ford and GM but somewhat lower for Toyota and Honda due to the exclusion of some of their plants⁶.

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 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. Measures and Hypotheses

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

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

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

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,

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

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. Figure 2 graphically summarizes our independent variables and their hypothesized relationship to DS .

Several covariates in X_{it} capture sales characteristics: sales trends and sales seasonality. 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. We include in X_{it} the following

⁸We 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.

two measures:

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

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

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 $STREND^+$ (because production capacity may lag the sales growth) and, DS is expected to be increasing in $STREND^-$ (because the firm may maintain a steady level of production even if the sales rate decreases). We include two measures to allow for different reactions to sales increases and decreases.⁹

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

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

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

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 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 definitions of these options are relatively standard, so it is possible to make comparisons of option

⁹Our 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.

intensity across models and years.¹⁰ Finally, $DEALERS_{it}$ is the number of dealerships in year t of model i 's brand. There are two reasons to suspect that demand fragmentation leads to higher DS . First, there may exist economies of scale in inventory management. 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. Second, demand fragmentation can lead to more variable demand, which can require more inventory to meet the same target service level, such as an in-stock service level (the probability of having a customer's preferred version in stock) or a fill-rate service level (the fraction of demand met immediately from stock) (see van Ryzin and Mahajan (1999)). Thus, we expect that DS is increasing in $DEALERS$ and $OPTIONS$, and decreasing in $SALES$.

In addition to demand fragmentation, $DEALERS$ and $OPTIONS$ can influence DS through other mechanisms. Adding options to a model may create additional production switching times in the assembly process, thereby reinforcing the positive relationship between $OPTIONS$ and DS . $DEALERS$ may influence inventory through a model loyalty mechanism. To explain, we use the term "model loyalty" to refer to the propensity of consumers to purchase a vehicle in a firm's assortment of a given model even if their most preferred version is not available immediately. If a model's loyalty is low, then a dealer needs to increase its target service level to ensure that sales are not lost due to poor availability (e.g., limited selection of colors and trim packages). On the other hand, if a model's loyalty is high, a dealer is unlikely to lose a sale even if inventory availability is low because consumers are then likely to substitute to another version or wait for their preferred version to be available. We hypothesize that $DEALERS$ is negatively associated with model loyalty - as the number of dealers for a brand increases, the dealers are more likely to be closer to each other and to dealerships from other brands, thereby increasing the choices available to consumers and lowering their model loyalty. Hence, $DEALERS$ is expected to be positively associated with DS . Note, the model loyalty mechanism relating $DEALERS$ to DS is consistent with the demand fragmentation mechanism that also relates the two of them.

Model loyalty is also likely to be influenced by the number of competing models in a segment - just as increasing the number of dealerships in an area gives consumers more choices, an increase in the number of models in a segment gives consumers more choice,

¹⁰Kekre 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.

thereby reducing model loyalty. Thus, included in X_{it} is $NMKT_{it}$, the number of models in the same segment as model i in year t . We expect DS is increasing in $NMKT_{it}$.

Competition can influence inventory through its impact on cost markups: if a product's margin decreases due to additional competition, then the firm has a lower incentive to offer a high service level. We are not able to observe $COSTMK_{it}$, the markup for model i in year t as a percentage of the model's cost, i.e., $(\text{price} - \text{cost}) / \text{cost}$, 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. Our cost markup estimation includes all non-truck vehicles sold in the U.S. in 1996-2004. As in Berry et al. (1995), we jointly estimate a model of consumer demand based on a random-coefficient multinomial logit and a reduced form supply equation to model marginal costs. On the consumer demand side, we include random coefficients for the following vehicle characteristics: price, size, acceleration, fuel efficiency, security and an indicator of market segment. We also include brand indicators and proxies for vehicle quality (obtained from Consumer Reports) as covariates without random coefficients. The supply side equation has the log of marginal cost as the dependent variable and covariates that include vehicle characteristics and variables describing the location where the vehicle is produced. To account for the endogeneity of price in the demand equation, we use the characteristics of other vehicles offered by the same firm and rival firms as instrumental variables.

The average estimated cost markup is 65%, which is line with the margins estimated by Berry et al. (1995) and other work using similar methodology (e.g. Petrin (2002)). To validate our cost markup estimates, we calculated the gross profit per vehicle for each make based on the estimated model markups and compared them to the dealerships' gross profit per vehicle published by Automotive News.¹¹ If dealerships get a fixed proportion of the supply chain profits, the estimated and actual profits should be proportional and highly correlated. The sample correlation between these measures is 80%. Further details on the cost markup estimation are shown in the online appendix.

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

¹¹The gross-profit data in Automotive News is collected by J.D. Power through a survey of U.S. dealers, which reports the average gross per vehicle for each brand.

model is manufactured in a plant that becomes more flexible, then we conjecture the plant will be able to produce in smaller batches, switch production more easily between models without substantial downtime periods and/or possibly more readily increase or decrease production by adding or subtracting shifts and/or overtime. As a result, a plant’s production should track sales more closely as it becomes more flexible. Therefore, we proxy production flexibility by the average absolute difference between production and sales, normalized by sales volume,

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

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).¹² Note, higher value of PS suggests a larger mismatch between production and sales. Nevertheless, we refer to PS as a proxy of production flexibility with the understanding that a large PS is associated with low flexibility. Hence we expect that DS is decreasing in PS .

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.¹³ 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. There is theory suggesting a positive effect of $UTIL$ on DS . For example, 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.¹⁴

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

¹²We 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 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.

¹³For 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.

¹⁴However, one could develop a model that predicts the opposite relationship. 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 .

components:

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

The random components δ_i and $\omega_{p(i)}$ represent time-invariant unobserved factors related to 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 loyalty, while factors in ϵ_{it}^m could include changes in model loyalty 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 $\epsilon_{it} = \epsilon_{it}^m + \epsilon_{p(i)t}^w$.

Table 1 describes the means of the variables used, grouped by manufacturer, and some other summary statistics for our 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.

6. Estimation Strategy and Results

There are several econometric concerns with the estimation of (1). Consequently, in this section we identify these concerns and report on four specifications designed to mitigate them.

A primary issue is that several of the factors included in (1) may be endogenous, i.e., controlled, at least in part, by the manufacturers. Consequently, 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). The inclusion of additional controls to the model can mitigate this endogeneity bias. We include in *all* four of our specifications the following control variables.

The regressions include year indicators to control for the random component τ_t . Make and segment controls are included to control for unobserved time-invariant differences in model loyalty across makes and segments (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. These controls may be important because a manufacturer may know that a particular segment has higher loyalty than others, thereby leading the manufacturer to choose a higher *COSTMK* for models in that segment, which implies a correlation between δ_i and *COSTMK*. To control for unobserved changes in model loyalty 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 produced¹⁵. To control for differences in replenishment leadtime (longer leadtimes should lead to higher *DS*), 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 . If model loyalty varies across models within a segment, these controls are useful for providing unbiased estimates of *COSTMK* and *SALES*. The concern with *COSTMK* has already been mentioned and the concern with *SALES* is similar: models with high loyalty may also have high sales, meaning that higher sales may be associated with lower *DS* rather than causal. This specification 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. *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.

Column (a) of Table 2 reports the estimates from our first specification. The signs of all the point estimates are consistent with theoretical predictions (except *STREND+*, which is

¹⁵Days-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.

positive and not significant), but not all the coefficients are different from zero with statistical significance. We defer our detailed discussion of the results from this specification to later in this section.

Our second specification is estimated without the model indicators. Here, we assume strict exogeneity, $E(u_{it}|X_{it}, W_{p(i)t}) = 0$, where X and W include all of the controls mentioned previously other than the model indicators; thus, we do not control for time-invariant unobservable differences across models (but the make and segment controls remain). 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 $E(\delta_i|X_{it}) = 0$ and $E(\delta_i|W_{p(i)t}) = 0$ 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. Note, it is also possible to include *SEASON* in this estimation.

Column (b) of Table 2 shows the estimates using RE. 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 $E(\delta_{it}|X_{it}) = 0$ and $E(\delta_{it}|W_{it}) = 0$ 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 loyalty: models with higher customer loyalty tend to have higher sales and markups, and higher loyalty allows the firm to choose lower inventory because consumers do not as readily switch to a competitor’s model when their most preferred version of the model is not available. Consequently, the empirical evidence suggests that controlling for model loyalty is important to get consistent estimates of the direct effect of cost markup and sales volume on inventory.

The third specification reintroduces the model controls and focuses on the estimation of *PS*. There is 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. However, we note that a mechanical relationship need not exist between the variability of inventory (*PS*) and the level of inventory (*DS*). 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, two products can have the same average inventory level but different inventory volatilities (batch sizes), or two products can have the same inventory volatility (batch size) but different average inventory levels (because they carry different safety stock levels). Nevertheless, to address this issue, we instrument PS using the following instrumental variables: the average PS of other models produced in the 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.

Column (c) of Table 2 reports results from the third specification. 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.¹⁶ The standard errors increase substantially for the estimated PS coefficient, 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 .

Our fourth specification deals with the issue of endogeneity in production planning choices and plant flexibility. If manufacturer's are aware of heterogeneity in production flexibility across their plants, they are likely to assign more platforms to the more flexible plants because those plants can better cope with the additional switching time each platform generates. Thus, there could be an association between $\omega_{p(i)}$ and $NPLATF$. To control for this (and possibly other) time-invariant unobservables in $\omega_{p(i)}$, this specification adds plant indicators (as well as model FE). Consequently, in this specification, both β and γ are estimated

¹⁶The sample excludes plants producing a single model and the first year in which a model is produced at a plant (which can be a new model or an existing model switching production between plants).

using variation across years only.¹⁷ In column (d) of 2 we report that 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=0.09). 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 loyalty. Specification (c), which corrects for a 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 now focus the analysis and discussion on the results provided by this specification.

To evaluate the economic significance of the results from our first specification (a), 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%. To provide another measure of the economic significance of the *DEALERS* and *PS* coefficients, Table 3 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 reported in the table 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. Furthermore, the potential inventory cost savings are substantial: nearly \$1billion for GM. Thus, our two main findings are:

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

1. increasing the number of dealerships, *DEALERS*, in a supply chain is associated with higher days-of-supply;
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.

Although the *DEALERS* effect is of large magnitude and significant, 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. (Recall, we control for differences across models, so *DEALERS* is not estimated with cross-sectional data.) Two makes did exhibit a considerable amount of variation in their dealership structure, 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.

While we identify an important effect regarding the dealership network on inventory, we are unable 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 to more demand fragmentation, which leads to higher demand variability and more inventory to achieve the same service level. Alternatively, more dealerships could lead to lower model loyalty, which leads to higher service levels. To distinguish between these two effects requires different data. Olivares and Cachon (2008) use cross sectional variation in inventory holdings of individual dealerships to provide evidence in support of the model loyalty effect - they find that the introduction of local competition causes incumbents to increase their inventory.

Aside from *DEALERS*, our other measures of demand fragmentation do not suggest a strong effect. For example, we did not find significant economies of scale associated with *SALES*, 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 which would lead to more inventory; but product differentiation offers a better match to heterogeneous customer preferences, making customers more loyal (less likely to substitute), which would lead to 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 find an important association between our proxy of production flexibility, *PS*, and our dependent variable, *DS*. *PS* measures inventory volatility and we suggest that as a plant becomes more flexible it generates less inventory volatility because it is better able to match its production to its demand. 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. Furthermore, *PS* appears to be capturing a measure of production flexibility beyond just the number of platforms produced at a plant, *NPLATF*, the aggregate scale of production, *SALES*, or measures of demand volatility, *SEASON*. 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 an increase in a model's cost markup (*COSTMK*) is 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 loyalty) can confound the direct effect of markups on inventory. Models with high loyalty have customers who do not substitute to a competitor's product when they do not find in stock a vehicle with their most preferred set of options, which allows the firm to carry less inventory relative to a model with lower loyalty. Furthermore, models with high loyalty tend to have higher markups, thereby suggesting that higher markups should be associated with lower inventory. Hence, a regression that does not include controls for model loyalty 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.

The results suggests that the effect of plant utilization, *UTIL*, is positive and significant. Two alternative explanations are consistent with this finding. An increase in utilization also increases the plant's 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.

Finally, we find a positive association between the number of models in a segment, *NMKT*, and *DS*. This is consistent with the theory that more substitutes within a segment make customer less prone to wait for a product that is out of stock. Consequently, stronger competition makes stock-outs more costly to a firm, leading to higher target service levels (and thereby higher inventories) to reduce the frequency of stock-outs.

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 production¹⁸. 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.

¹⁸These include COROLLA after year 2001, and all the model-years of ACCORD, CAMRY and MAXIMA; a total of 15 observations.

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 loyalty. 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 econometric model includes $NMKT$ – the number of models offered in a segment by all companies in the market – as a factor that influences model loyalty. However, it may be possible for this effect to be different depending on whether the substitutes of a model are offered by the same company or by a rival company. To explain, a company that offers a more attractive assortment (more models to choose from) may not need to carry as much inventory to prevent losing customers to a competitor, because customers are more likely to purchase within the company’s assortment when facing a stockout. We tested this by replacing $NMKT_{it}$ with two measures: $NMKT_OWN_{it}$, which counts the number of models offered in the segment by the same company offering model i , and $NMKT_OTHER_{it}$, which includes the number of models offered by all other companies. The coefficient of $NMKT_OTHER$ is similar to that of $NMKT$ (0.247 with standard error 0.086) and the coefficient on $NMKT_OWN$ is -0.03 and not statistically significant. Hence, DS is increasing in the number of substitute models offered by rival companies, but is relatively insensitive to the number substitutes offered by the same company. This provides further evidence that $NMKT$ is capturing the effect of competition and model loyalty.

Our results suggest that including model FE is useful to control for unobserved model loyalty to get consistent estimates of the effect of $COSTMK$ on inventory performance. But if model loyalty changes across time, then model FE does not control completely for this confounding effect. To test this, a proxy that captures longitudinal variation in model loyalty 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 highest rating¹⁹. 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 smaller²⁰. 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. If the Consumer Reports ratings are a reasonable proxy for changes in a model’s customer loyalty, this analysis suggests that unobservable variation in model loyalty over time is not a major confounder in our results.

Specification (d) in Table 2 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).²¹ 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).

Our measure of production flexibility, *PS*, is strongly correlated with measures of demand variability: the correlation between *PS* and the coefficient of variation of sales is approximately 0.55. Hence, it may be possible that *PS* is picking up the effect of promotion activity (which is related to demand variability) in addition to production flexibility. To test this, we included in the regression a measure of demand variability– the coefficient of variation of the model sales during the year, *CVSALES_{it}*. The coefficient on *PS* remains similar in sign and magnitude (point estimate is 0.241 with standard error 0.029), and the coefficient of *CVSALES_{it}* is -0.11 and statistically significant. The coefficients on all the other covariates

¹⁹Consumer Report classification of vehicles includes more segments than the four we use, decomposing the car and SUV segments into multiple groups (luxury, middle/large size, etc).

²⁰The sample of Consumer Reports models tends to include higher selling vehicles than in the base sample (122 versus 104 thousand vehicles).

²¹We 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.

do not change much. While demand variability seems to have some effect, it is relatively small compared to the other factors considered in this study. More importantly, the effect of PS remains large and significant after controlling for $CVSALES_{it}$, suggesting that PS is capturing mismatches between production and sales rather than just demand variability.

The specifications in Table 3 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 $NPLATF$.

The results in Table 2 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 product variety. A new measure was defined based on Ward’s platform classification, which is different from Harbour’s platform definitions.²² 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).

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

²²Ward’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.

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 is costly 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 20th 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, manufacturers find it difficult to 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)). GM and Chrysler have recently taken steps to drastically reduce the number of dealerships in their network, but only in the context of threatened and actual bankruptcy proceedings (Vlasic and Bunley (2009)).

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 US 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.
- Olivares, Marcelo, Gerard Cachon. 2008. Competing retailers and inventory: An empirical investigation of GM dealerships in isolated markets. *Management Science* (*forthcoming*) .
- Petrin, A. 2002. Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy* **110**(4) 705–729.
- 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.
- Vlasic, Bill, Nick Bunley. 2009. Chrysler Plans to Shut 1 in 4 of Its U.S. Dealers. *New York Times* .
- 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

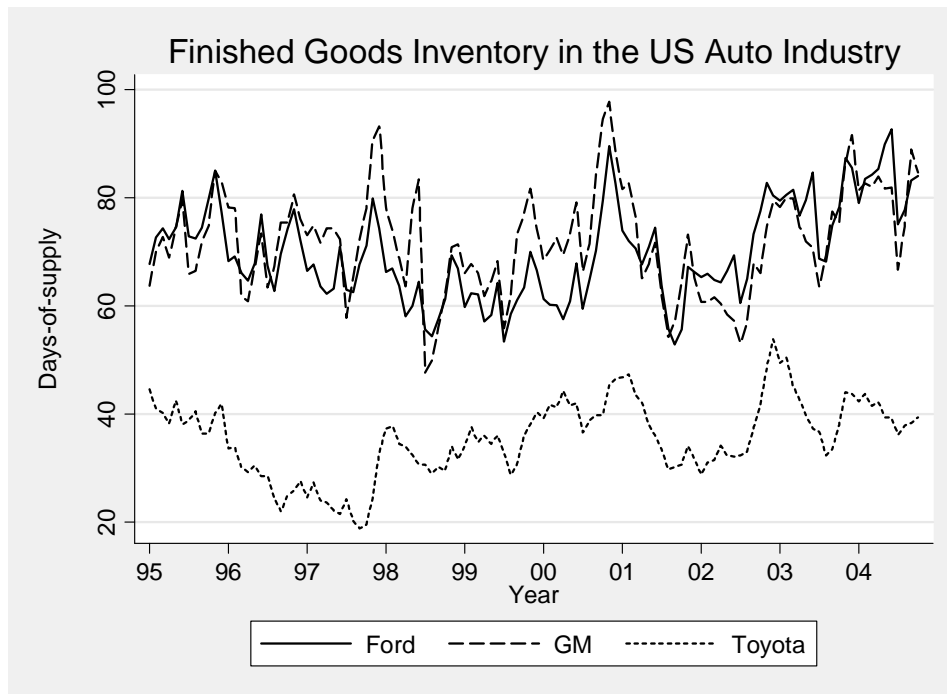


Figure 1 – Inventory of three auto manufacturers. Days-of-supply is calculated as the aggregate inventory at the end of each month divided by the average daily sales rate in the following two months. Inventory includes all finished vehicles in US territory, including inventory in the plant, in ports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

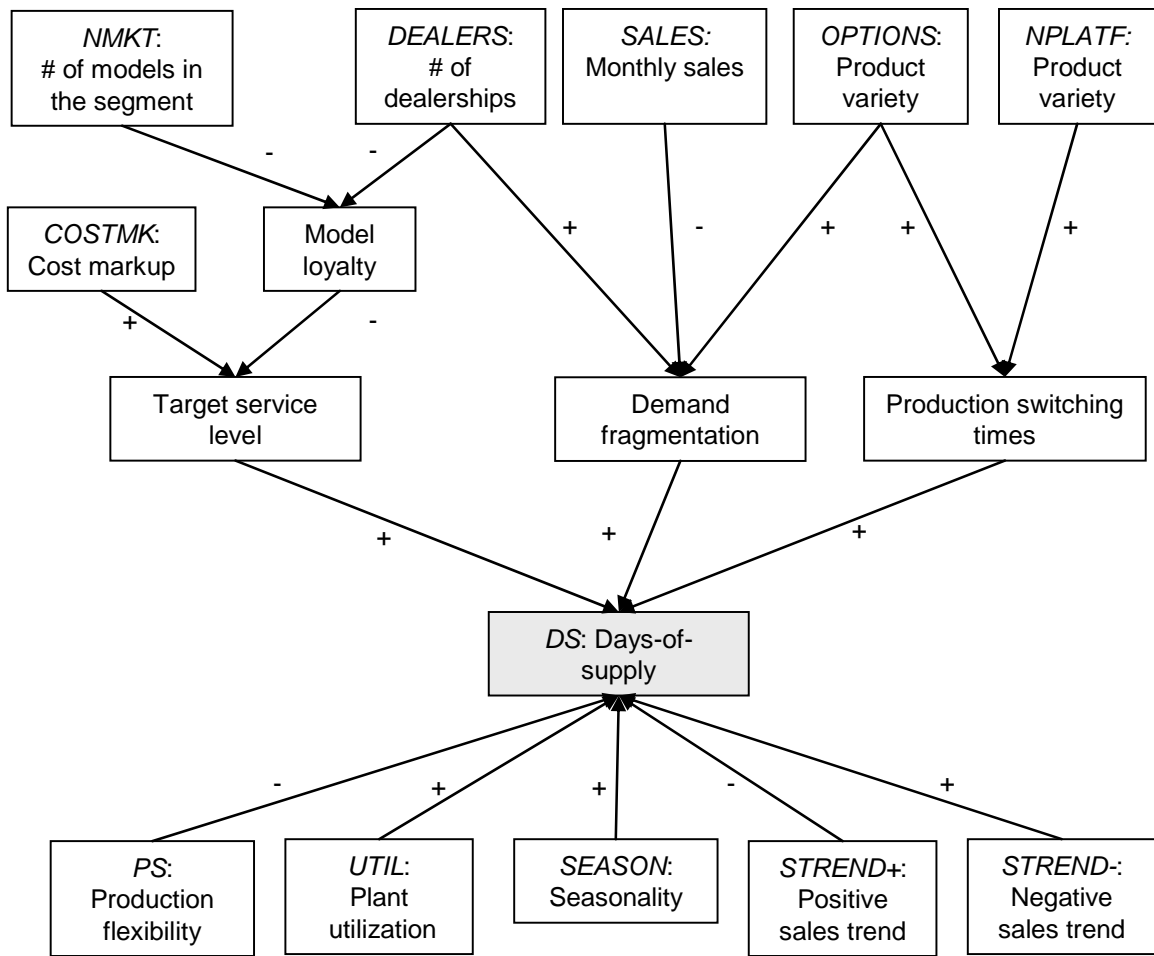


Figure 2 – Theoretical factors influencing days-of-supply. Covariates included in our model are indicated with italics. Signs indicate the hypothesized relationship, where “-” indicates an inverse relationship and “+” indicates a positive relationship.

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

Model: $DS_{it} = \beta X_{it} + \gamma W_p(0)_t + u_{it}$

	(a)	(b)	(c)	(d)
SALES	-0.016 (0.040)	-0.105** (0.024)	-0.001 (0.047)	-0.054 (0.045)
STREND-	0.019** (0.006)	0.016* (0.006)	-0.002 (0.007)	0.019** (0.006)
STREND+	0.005 (0.006)	0.002 (0.006)	-0.015* (0.007)	0.006 (0.006)
DEALERS	0.528* (0.207)	0.483* (0.210)	0.699** (0.237)	0.452* (0.218)
OPTIONS	0.058* (0.024)	0.072** (0.021)	0.016 (0.026)	0.057* (0.025)
COSTMK	0.439** (0.154)	0.191^ (0.110)	0.311^ (0.171)	0.554** (0.162)
NMKT	0.220** (0.071)	0.213** (0.065)	0.266** (0.086)	0.189* (0.079)
UTIL	0.178** (0.041)	0.129** (0.035)	0.211** (0.046)	0.226** (0.045)
NPLATF	0.047 (0.042)	0.073* (0.033)	0.047 (0.054)	0.088^ (0.052)
PS	0.194** (0.027)	0.218** (0.026)	0.213^ (0.111)	0.175** (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 2 – 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. Overall R-square is reported in column (b), while within R-square is reported in the others. 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.

Manuf.	Days of Supply	<u>% reduction in days-of-supply</u>		<u>Adjusted Days of Supply</u>		Inv. Cost Reduction M\$
		PS	DEALERS	Estimate	95% CI	
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 3 - 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.