Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Does Adding Inventory Increase Sales? Evidence of a Scarcity Effect in U.S. Automobile Dealerships

Gérard P. Cachon, Santiago Gallino, Marcelo Olivares

To cite this article:
Gérard P. Cachon, Santiago Gallino, Marcelo Olivares (2018) Does Adding Inventory Increase Sales? Evidence of a Scarcity Effect in U.S. Automobile Dealerships. Management Science

Published online in Articles in Advance 12 Jul 2018

. https://doi.org/10.1287/mnsc.2017.3014

informs
http://pubsonline.informs.org/journal/mnsc/

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–17
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Does Adding Inventory Increase Sales? Evidence of a Scarcity Effect in U.S. Automobile Dealerships

Gérard P. Cachon,[a] **Santiago Gallino,**[b] **Marcelo Olivares**[c]

[a] The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104; [b] Tuck School of Business, Dartmouth College, Hanover, New Hampshire 03755; [c] Industrial Engineering, Universidad de Chile, Santiago, Chile

Contact: cachon@wharton.upenn.edu, http://orcid.org/0000-0002-7046-8835 (GPC); santiago.gallino@dartmouth.edu, http://orcid.org/0000-0002-8906-2691 (SG); molivares@u.uchile.cl, http://orcid.org/0000-0001-8465-624X (MO)

Received: September 16, 2016
Revised: June 9, 2017; November 2, 2017
Accepted: November 16, 2017
Published Online in Articles in Advance: July 12, 2018

https://doi.org/10.1287/mnsc.2017.3014

Abstract. What is the relationship between inventory and sales? Clearly, inventory could increase sales: expanding inventory creates more choice (options, colors, etc.) and might signal a popular/desirable product. Or, inventory might encourage a consumer to continue her search (e.g., on the theory that she can return if nothing better is found), thereby decreasing sales (a scarcity effect). We seek to identify these effects in U.S. automobile sales. Our primary research challenge is the endogenous relationship between inventory and sales—e.g., dealers influence their inventory in anticipation of demand. Hence, our estimation strategy relies on weather shocks at upstream production facilities to create exogenous variation in downstream dealership inventory. We find that the impact of adding a vehicle of a particular model to a dealer's lot depends on which cars the dealer already has. If the added vehicle expands the available set of submodels (e.g., adding a four-door among a set that is exclusively two-door), then sales increase. But if the added vehicle is of the same submodel as an existing vehicle, then sales actually decrease. Hence, expanding variety across submodels should be the first priority when adding inventory—adding inventory within a submodel is actually detrimental. In fact, given how vehicles were allocated to dealerships in practice, we find that adding inventory actually lowered sales. However, our data indicate that there could be a substantial benefit from the implementation of a "maximize variety, minimize duplication" allocation strategy: sales increase by 4.4% without changing the total number of vehicles at each dealership.

History: Accepted by Vishal Gaur, operations management.
Funding: M. Olivares acknowledges funding from the Chilean National Fund for Scientific and Technological Development (Conicyt) [Grant FONDECYT-1161333] and the Complex Engineering Systems Institute [PIA-FB0816].

Keywords: retail operations • inventory competition • empirical operation management • instrumental variables • automobile industry • scarcity

1. Introduction

In early 2008, before the financial crisis, car dealerships in the United States held enough vehicles to cover sales for 75 days (WardsAuto market data). However, immediately following the financial crisis automakers began drastic reductions in their inventories. By January 2010, days-of-supply for the industry had dropped to less than 49, leading many dealers to complain that their low inventories were negatively affecting sales (Barkholz 2010). Were those complaints justified?

Adding inventory can increase sales for several reasons. For example, additional inventory reduces the chance of running out of stock (which constrains sales) and allows a dealer to expand the variety of options (e.g., trim, colors, options) to a customer. This increases the odds that a customer finds a vehicle that sufficiently matches her preferences, thereby increasing the likelihood of a sale. Inventory can also change preferences. For example, seeing many cars on a dealer's lot might cause a customer to infer that the car is popular (a dealer carries many cars only if the model is popular), thereby making the car more desirable to the customer.

There are also downsides to ample inventory. For instance, if there are many cars on a dealer's lot, then a customer might infer that the car is not popular, and that it must not be popular for a reason, so the customer becomes less likely to purchase. Or, seeing that many units are available, a customer may become more likely to continue her shopping/search process at other dealers because she believes that if she decides to return to the dealership, the car she desires will still be available. Once the customer leaves, she might not return (because there is a chance she will find a better alternative elsewhere or she might change her mind about even purchasing a vehicle), so this more active search lowers the dealership's sales.

In general, for simplicity, we use the label "variety effect" for any mechanism that assigns a positive relationship between inventory and demand, with the understanding that some mechanisms may not be directly related to product variety (as in when multiple units of the same product signals popularity/desirability). And we use "scarcity effect" for any mechanism with a negative relationship. Our objective is to empirically evaluate the strength of these effects in the U.S. auto industry and then to use those estimates to recommend how existing vehicles should be allocated across dealerships to maximize sales.

While it is possible to identify several mechanisms that connect inventory to sales, estimating the relationship between inventory and sales is complex primarily because it is reasonable to believe that inventory is chosen endogenously. For example, a simple plot reveals a positive relationship between the amount of inventory a dealer carries and the dealer's average weekly sales. But dealers that operate in larger markets are expected to carry more inventory and have higher sales even if inventory has no influence on demand merely because a firm rationally needs to carry more inventory when it serves more demand. To overcome this selection effect, we estimate the influence of inventory using only observed variation within dealer–model pairs rather than variation across dealerships and models. This approach is valid given the assumption that a dealer's market conditions are reasonably constant in our six-month study period (e.g., there is little change in local factors like demographics, population, or the degree of competition the dealer faces). However, even within a dealer–model pair, there is a concern that a dealer may change her inventory level in anticipation of changes in demand. For example, the dealer may build inventory because of a planned promotion. In that situation it is incorrect to conclude that the larger inventory caused the higher sales. To overcome this issue, we exploit shocks to dealers' inventories due to extreme weather events at upstream production locations. Extreme weather may disrupt production via a number of possible mechanisms (e.g., delays in inbound or outbound shipments, worker absenteeism, adjustments in production schedules in anticipation of weather or in response to weather, etc.) and also is independent of dealer demand (as production generally occurs at a considerable distance from the dealership), thereby providing a valid instrument that allows us to estimate the causal impact of inventory on sales.

2. How Inventory Impacts Sales

In this section we describe several mechanisms by which more inventory can increase sales and then discuss mechanisms that predict the opposite.

At a basic level, it is intuitive that more inventory can help to avoid running out of stock, thereby increasing sales. For example, if q units of inventory are available of a single item and its stochastic demand is d, then expected sales of that item, $E[\min(q,d)]$, is increasing in q. Related to this "stockout effect," increasing inventory can increase the variety of available products, thereby increasing the choice available to customers. Expanding the set of available choices increases sales because customers are more likely to find an item that suits their preferences (see, e.g., Train 2009, Talluri and van Ryzin 2004, Smith and Achabal 1998). Kalyanam et al. (2007) provide empirical evidence for this in a nonfashion apparel category. However, they are unable with their data to directly measure the relationship between inventory and sales.

Inventory can influence a consumer's engagement in the purchasing process. For example, if a consumer is not aware of an item, the consumer cannot even consider purchasing it—a large inventory may act like a billboard and increase awareness, which increases sales. Or, a consumer may infer that a large inventory implies a low price (e.g., the item must be on promotion or the dealer will be willing to negotiate a good deal), thereby motivating the consumer to include the item in her consideration set (see Zettelmeyer et al. 2006 for a study on the effect of dealership inventory on prices). If search is costly, then consumers are more likely to visit (and therefore buy from) a dealer that has a reputation for higher inventory—nobody likes to go to a store only to discover that the desired item is unavailable (e.g., Deneckere and Peck 1995, Dana and Petruzzi 2001, Bernstein and Federgruen 2004, Su and Zhang 2009, Matsa 2011).

It is also possible that inventory changes preferences. For example, a consumer might infer from a large inventory that the item has good quality (why else would the dealer have so many), thereby making the item more desirable to the consumer—a good-quality item has useful features and durability. For example, Balakrishnan et al. (2004) assume that inventory increases demand and study how this effect influences single-product inventory decisions in a deterministic environment.

In contrast to the various mechanism through which inventory potentially increases sales, there are several mechanisms that lead to a scarcity effect in which more inventory actually lowers sales. This could happen if consumers infer that an item with ample inventory is unpopular or of low quality—there must be many units because nobody is buying the item (e.g., Balachander et al. 2009, Stock and Balachander 2005). Or, a consumer might prefer an item that is perceived to be exclusive or rare, as in a collectible (e.g., Brock 1968, Brehm and Brehm 1981, Worchel et al. 1975, Tereyağoğlu and Veeraraghavan 2012). This may apply to some specialty vehicles in the auto industry, but probably not to the sample of mainstream vehicles we consider.

If it is costly for consumers to consider all possible options, then low inventory may imply a low variety of options and higher confidence that a good option has been identified (e.g., Kuksov and Villas-Boas 2010). Similarly, high inventory and high variety may create confusion or frustration (too many options to know where to begin), thereby leading to lower demand and sales (e.g., Iyengar and Lepper 2000, Schwartz 2004, Gourville and Soman 2005).

A large inventory may indicate that a product will be available later on at a good price (because the dealer may need to discount the item), thereby encouraging consumers to wait before buying, which lowers current sales (e.g., Aviv and Pazgal 2008, Su and Zhang 2008, Cachon and Swinney 2009). In contrast, with a low current inventory consumers not only anticipate that the price will not fall, but also that the item may not be available in the future. This can lead to a "buying frenzy" in which the low current inventory creates a sense of urgency among consumers to buy immediately (DeGraba 1995, Qian and van Ryzin 2008). Soysal and Krishnamurthi (2012) use data from a specialty apparel retailer to document empirically that consumers do wait strategically for end-of-season discounts. Nevertheless, they find via a counterfactual analysis that reducing inventory reduces sales (through a stockout effect) even if profits may increase (because more items are sold at the full price).

Inventory may influence consumer search. Say a consumer finds a vehicle that she likes at a dealership. If the dealer has only one of that type of car, she may be inclined to stop her search and just buy the car—if she continues her shopping at other dealers, then she risks not finding a better car and losing the current car to another customer. But if the dealer has several of her desired cars, she may be more inclined to continue her search, and that search may lead her to make a purchase from some other dealership. (See Cachon et al. 2008 for a model in which variety influences the degree of consumer search.) Alternatively, ample inventory may encourage her to delay making a purchase commitment so that she might further mull over the decision, or even cause her to forgo the purchase entirely.

To summarize, there are several mechanisms that suggest more inventory increases sales (ample inventory enables a better preference match, increases awareness, signals popularity, indicates availability, and suggests the potential to obtain a good price). For simplicity, we collectively refer to these as variety effects given that variety is likely to be a key factor in consumer purchasing decisions in the auto industry. In contrast, other mechanisms suggest more inventory decreases sales (ample inventory reduces the urgency to purchase immediately while encouraging additional search, signals an unpopular vehicle, creates overwhelming choice, and suggests that prices will soon be lowered). We refer to these as scarcity effects. We seek to measure these effects in the auto industry using fine-grained field data. We find that the way in which inventory is added to a dealer matters considerably: Does the additional inventory expand the available set of submodels or increase the number of units available within a submodel? This knowledge allows us to devise a new method for allocating existing vehicles that yields higher sales without changing which vehicles are produced.

3. Data Description and Definition of Variables

As a general reference, during the period of our study, six car companies accounted for about 90% of sales in the U.S. auto market. The company we focus on, General Motors (GM), captured 25% of the market. This market share was distributed across several different brands: Chevrolet, GMC, Pontiac, Buick, Saturn, Cadillac, and Hummer.

The data used in our analysis can be separated into two groups. The first group includes the inventory and sales information for the dealers in our sample. The second group includes geographic location, weather information for all of the GM dealers in our sample and all GM plants located in the United States and Canada.

3.1. Dealer's Sales and Inventory Data

We obtained, via a web crawler, daily inventory and sales data from a website offered by GM that enables customers to search new-vehicle inventory at local dealerships. The data collection was done from August 15, 2006, to February 15, 2007, and includes a total of 1,289 dealers in the following states: California, Colorado, Florida, Maine, Nebraska, Texas, and Wisconsin. These states are geographically dispersed and somewhat geographically isolated—they may border with Mexico or Canada or have a substantial coastline. The dealers in the sample are all of the GM dealers in those states, and they represent approximately 10% of all GM dealers in the United States for the period under analysis.

The crawler collected specific information for each vehicle at a dealer's lot, such as its trim level, options, list price, and vehicle identification number (VIN). Our sample of GM vehicles includes all cars and a large portion of light-truck models manufactured and sold in the United States and Canada. VINs uniquely identify all vehicles in the United States. Thus, they provide three key pieces of information. First, the VINs allow us to identify when a new car arrived at a dealer and when a sale happened (a vehicle is removed from a dealer's inventory). Second, the VIN code identifies the particular plant where the vehicle was produced even if the model is manufactured at multiple plants. Finally, the VINs provide us with information regarding dealer transfers—we can observe when a vehicle is removed from one dealer's inventory and added to another dealer's inventory within the state. If a vehicle leaves a dealer in week t and does not reappear in another dealer's inventory in week $t+1$, then we

Figure 1. Dealer and Plant Locations in Our Sample

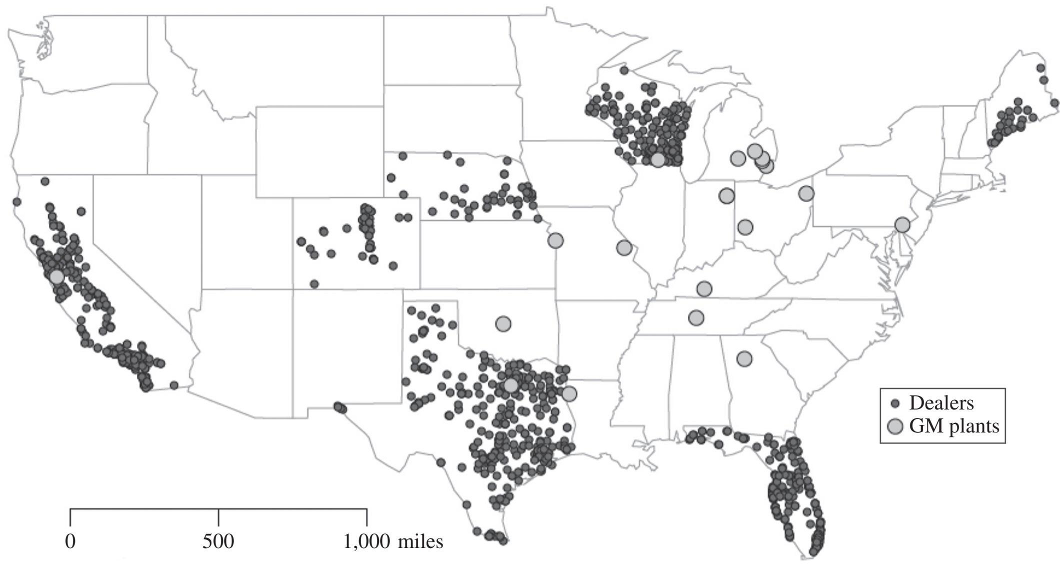

code this as a sale. Otherwise, it is coded as a transfer. For example, if car A is transferred from dealer 1 to dealer 2 and then sold at dealer 2, a sale is counted only at dealer 2. We can only observe transfers between dealerships within the same state. We anticipate that we observe the majority of transfers because transfers occur in a limited geographic area.[1] We removed from our sample a limited number of dealerships that opened or closed during the period under analysis.

Figure 1 shows the geographic location of GM plants and dealers, and Table 1 summarizes the number of dealers in each state in our sample.

3.2. Geographic Location and Weather Data

For each dealer and all 22 GM plants supplying vehicles in our sample (located in the United States and Canada), we obtained their address and exact geographic location (longitude and latitude) from GM's website.

Our first source of weather data was the National Weather Service Forecast Office (NWSFO). We obtained daily weather information for every dealership and plant location in our sample for the period August 15, 2006, to February 15, 2007. Included in the sample are each day's maximum, mean, and minimum values for the following weather variables: temperature, wind speed, humidity, pressure, visibility, and dew point. We also obtained information on the type of event during a day (rain, thunderstorm, snow, etc.). We identified the closest weather station to each plant and each dealer. The selected weather stations are close to our plants with a mean and median distance of 12 and 10 miles, respectively. No plant is further than 32 miles from its corresponding weather station. To assess whether a station's weather is likely to be similar to the weather at its nearby plant, we constructed a sample of weather stations that are between 30 and 60 miles apart. In this sample, the correlation in our weather variables is no less than 95%, suggesting that the weather reported at the nearby weather station is representative of the weather at the plant.[2]

The second source of data was the National Climatic Data Center of the National Oceanic and Atmospheric Administration (NOAA). On their website they share access to the "Storm Events Database" (NOAA 2015). The type of storms and events recorded on this database are consistent with the Storm Data listed in table 2 of section 9 of National Weather Service (NWS) Directive 10-1605 (NOAA 2016). The storm data are reported at the FIPS level. The county FIPS number is a unique number assigned to the county by the National Institute of Standards and Technology (NIST) or NWS forecast zone number. We matched production plants to county and zip codes to establish a one-to-one link between the Storm Events Database and the production plants.

Table 1. Dealers by State in our Sample

State	Number of dealers
California	355
Colorado	67
Florida	237
Maine	31
Nebraska	50
Texas	366
Wisconsin	183
Total	1,289

4. Model Specification

We seek to estimate the impact of inventory and variety on sales. The available data were used to construct

Table 2. Weather Variables Included in the Study

Variable	Description	Avg.	SD	Median	Min	Max
Wind	Number of days in which a wind advisory is issued by the National Weather Service Forecast Office	0.01	0.11	0.00	0.00	1.00
Cloud	Average cloud cover during the week (0 = no clouds; 8 = sky completely covered)	3.83	1.98	3.71	0.00	7.85
Fog 1	Weeks with one day with fog during the week	0.15	0.36	0.00	0.00	1.00
Fog 2–3	Weeks with two or three days of fog during the week	0.09	0.29	0.00	0.00	1.00
Fog 4–7	Weeks with more than three days of fog during the week	0.02	0.13	0.00	0.00	1.00
Rain 1–2	Weeks with one or two days of rain during the week	0.37	0.48	0.00	0.00	1.00
Rain 3–5	Weeks with three to five days of rain during the week	0.33	0.47	0.00	0.00	1.00
Rain 6–7	Weeks with more than five days of rain during the week	0.01	0.09	0.00	0.00	1.00
Snow 1	Weeks with one day of snow during the week	0.12	0.32	0.00	0.00	1.00
Snow 2–4	Weeks with two to four days of rain during the week	0.11	0.31	0.00	0.00	1.00
Snow 5–7	Weeks with more than four days of rain during the week	0.03	0.16	0.00	0.00	1.00
High Temp 1	Weeks with one day of high temperature, above 90 degrees Fahrenheit, during the week	0.12	0.32	0.00	0.00	1.00
High Temp 2–5	Weeks with two to five days of high temperature, above 90 degrees Fahrenheit, during the week	0.07	0.25	0.00	0.00	1.00
High Temp 6–7	Weeks with more than five days of high temperature, above 90 degrees Fahrenheit, during the week	0.02	0.13	0.00	0.00	1.00

a panel data set where the unit of analysis is the log of sales plus one of a particular vehicle model i at a specific dealership j during a week t ($Sales_{ijt}$). Expected sales during a week are influenced by the total number of vehicles available at the dealership during the week ($Inventory_{ijt}$), the number of varieties of the model that where available ($Variety_{ijt}$, to be described in more detail shortly), plus other factors that could influence the demand for vehicles at the dealership. Figure 2 illustrates the relationship between the key variables in our analysis—sales, inventory, and variety.

Figure 2 shows multiple effects between the three key variables. First, there is a direct effect of inventory on sales. An example of this effect is when a low level of inventory signals low future availability of the vehicle model and leads to a "buying frenzy" behavior, or when a high level of inventory signals lower prices and therefore increases sales. Therefore, the sign of this relationship is ambiguous. Second, there is a direct

Figure 2. Relationship Between Sales, Inventory, and Variety

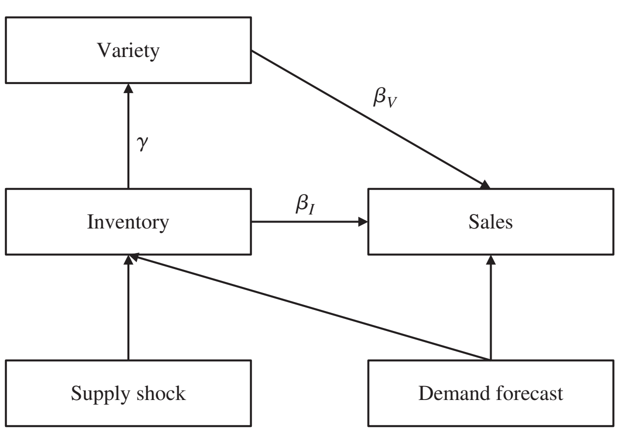

effect of variety on sales, as when more variety leads to a better match of customer preferences, thereby increasing sales. Higher variety could also lead to more confusion in choosing among too many options, lowering sales. Hence, the sign of this relationship is also ambiguous. Third, there is an indirect effect of inventory on sales through variety: adding inventory can lead to an increase in variety, which in turn could affect sales.

Using the indexes i for dealership, j for model, and t for week, the following regression equation is used to estimate the impact of inventory and variety on sales:

$$Sales_{ijt} = \beta_0 + \beta_V Variety_{ijt} + \beta_I Inventory_{ijt} + \text{Controls} + \varepsilon_{ijt}. \tag{1}$$

The error term ε represents factors that affect sales that are unobservable in the data. Dealerships and manufacturers may predict some of these factors in advanced and use them in their demand forecast to choose inventory levels and model variety requirements (see Figure 2 for an illustration). Hence, ε is likely to be positively correlated with *Inventory* and *Variety*, making these variables endogenous in Equation (1).

To obtain unbiased coefficients for β_V and β_I, we need to find a set of variables that can be excluded from the sales equation (i.e., do not affect sales directly) but have a direct impact on inventory and variety.

If weather at the plant affects its production, then weather shocks at the plant affect the inventory level at the dealerships. Because most of the plants are located far away from the dealerships in our study, weather shocks at the plants should be unrelated to the local demand for autos. Weather shocks at the plants, defined by the vector of covariates W, are excluded from Equation (1) while affecting inventory levels at the dealer and therefore are valid instrumental variables for *Inventory*.

Although dealerships can control the number of vehicles of a particular model that they receive, they typically have little control on the exact submodels that are allocated to them. The variations in variety after controlling for inventory levels should be for the most part unrelated with the demand forecasts or other unobservable factors related to demand. By explicitly controlling for the inventory level, we are able to identify separately the effect of variety. However, we cannot completely rule out the possibility that *Variety* is also endogenous in Equation (1). Hence, *Variety* has to be instrumented to be certain that consistent estimates are obtained for the coefficients in Equation (1).

We instrument $Variety_{ijt}$ using a "Hausman-type" instrument defined by the average variety for model j at the three closest dealers (h) to dealer i that are at least 150 miles away from dealer i, which we labeled $NearVariety_{hjt}$. Distant dealers are selected to avoid any potential correlation between their variety and the sales at dealer i. However, we can verify that the variety for a particular model j within a region is correlated across dealers because dealers within a region tend to receive vehicles from a common production mix through the same distribution channels. Hence, variety at the other dealers in the region, but outside a dealer's market, are excluded from Equation (1) while affecting variety levels at the dealer. Consequently, the explanatory variable $NearVariety_{hjt}$ is a valid instrumental variable for $Variety_{ijt}$ in Equation (1).

Given the valid instruments for the two endogenous variables (*Inventory* and *Variety*), we are able to obtain unbiased coefficients for β_I and β_V by implementing a two-stage least squares (2SLS) estimation approach. Although so far we emphasize the use of instrumental variables to address the endogeneity bias that exists between inventory, variety, and sales, the instrumental variable approach also allow us to address any concern arising from other potential omitted variables in the model.

4.1. Controls

Equation (1) includes several control variables. Model–dealership fixed-effects are included to control for invariant characteristics of each dealer and model: dealer location, the average popularity of a model at a particular dealership, the intensity of competition a model faces at each dealer, the manufacturer's suggested retail price (MSRP) of a model, and the average discount policy a dealer offers for a particular model, among other characteristics. Seasonal controls include dummy variable to account for changes in sales across weeks. This is implemented by grouping dealers into four geographic regions: {Florida, Texas}, {Colorado, Nebraska}, {Maine, Wisconsin}, and {California}. Let $r(i)$ be the region containing dealership i. We include the set of dummy variables $Seasonal_{r(i)t}$ to control for different seasonal patterns across geographic regions—e.g., a different weekly sales pattern in Texas than in Wisconsin. In addition, the regression includes covariates capturing local weather at each dealership to control for the effect of local weather on sales and demand forecasts. Steele (1951) and Murray et al. (2010) present examples of how local weather affects retail sales, and Busse et al. (2014) show evidence of this effect specifically for car dealers. This control is particularly important to ensure that plant weather is an exogenous instrumental variable because, for some dealerships, local weather could be correlated with plant weather.

4.2. Measuring Variety

To identify which of the main effects of inventory on sales described earlier dominates, we identify separately the impact of our two measures of availability—inventory and variety. For example, a negative effect of $Variety_{ijt}$ would suggest that the confusion effect dominates the impact on sales. Although $Inventory_{ijt}$ can be objectively defined as the number of vehicles available for a model, variety could be defined in different ways depending on the relevant product characteristics that are considered by customers when making their purchase decision. For example, a customer wanting to buy a Chevrolet Malibu may consider two vehicles with different horsepower as two different products, but could be indifferent on the color of the car. To measure *Variety*, it is necessary to define a set of attributes that describes relevant differences across vehicle options within a model. See Hoch et al. (1999) for a framework on how customers perceive variety.

The VIN of a vehicle contains information about vehicle characteristics, including the model, body style, engine type, and restraint type. We use all of these relevant characteristics reported in the VIN to define the different possible variants of a model, and we refer to each variant as a *submodel*. The variable $AvailVar_{ijt}$ is the number of submodels of a model j available at dealership i during week t. The assumption here is that the variety information included in the VIN describes relevant differences across vehicle options from the customer perspective.

Table 3 summarizes the number of different submodels observed in our data and the average $Variety_{ijt}$ observed at the dealerships for a sample of models. The table reveals that there is variation in the number of submodels available across the set of models. Hence, it is plausible that the impact of variety is different across models: for example, adding one more submodel of a Cobalt (which has many submodels) can have a smaller impact than adding one more submodel of an Equinox (which has fewer available submodels). To account for this, the amount of available variety can be measured

Table 3. Model Variety for the Top 10 Selling Models

	Total model variety (*MarketVar*)	Average variety available (*AvailVar*)
Cobalt	18	3.5
Equinox	4	2.2
G6	37	6.1
HHR	4	2.9
Impala	10	3.7
Suburban	18	4.5
Tahoe	13	4.0
TrailBlazer	10	2.1
Saturn VUE	5	4.6
Yukon	30	8.6
Average	14.9	4.2

Notes. MarketVar is the maximum number of variants that could be produced for the model. *AvailVar* is the number of variants with at least one unit during a particular week.

relative to the number of submodels that exist for that model. Denote $MarketVar_j$ as the total number of submodels that were observed for model j. Our measure of variety is defined as

$$Variety_{ijt} = \frac{AvailVar_{ijt}}{MarketVar_j}. \quad (2)$$

Tables 4 and 5 present descriptive statistics and correlations, respectively, for our main variables.

Our period of analysis comprises the production period of model-year 2007 (most of the model changeovers occurred between June and August). However, there are some model-year 2006 vehicles observed in the inventory (less than 4.5%) that are considered as different submodels (of the corresponding model) and therefore add to variety. These vehicles remaining from model-year 2006 are typically discounted through rebates and dealer incentives. Hence, a positive correlation between *Variety* and *Sales* could be driven by seasonal discounts of the 2006 model-year vehicles. This issue is addressed in two ways. First, the control variables include week fixed effects, which would capture the seasonal discounts following the model changeover period. Second, *Variety* is instrumented with the variety of other dealerships. Hence, the effect of variety is identified using the variation of inventory across different regions, controlling for seasonal patterns that affect variety over time (such as price discounts in model changeover periods).

Table 4. Descriptive Statistics

Variable	Mean	SD	Min	Max
Sales (units)	0.75	1.47	0	35
Inventory (units)	11.67	14.83	0	193
Variety (count)	3.14	2.664	0	21
Variety Percentage (%)	36.76	27.35	0	100
Total number of observations: 216,438				

Table 5. Correlation Table

	Sales	Inventory	Variety	Variety Percentage
Sales	1.00			
Inventory	0.64	1.00		
Variety	0.39	0.70	1.00	
Variety Percentage	0.23	0.31	0.12	1.00
Total number of observations: 216,438				

4.3. Weather Instrumental Variables

Our instrument can work in several ways: bad weather can affect the supply of parts to the production line, slowing the production process; weather conditions can affect employee behavior both in their task performance and by increasing absenteeism; or weather can delay shipments of vehicles to dealers. Alternatively, production schedules could be altered in anticipation of weather.

To use weather as an instrument, we considered the data obtained from NOAA on "storm events." The benefit of considering this variable resides in its simplicity and transparency since the decision of whether there is an extreme event or not is define by NOAA. This definition is location specific: NOAA uses different criteria to announce the extreme event at different locations (e.g., two inches of snow in North Carolina can be characterized as an extreme event, whereas the same snowfall in Massachusetts would not). To implement the analysis, we group the different events reported by NOAA into six variables according to the type of event: thunder storm, winter storm, tropical storm, heat wave, high wind, and flood. Each one of these six variables is an indicator of whether a particular production plant was affected by an extreme weather event during week t.

4.4. An Alternative Estimation Approach of the Overall Effect of Inventory

Assuming that inventory has a linear effect on variety as characterized by the γ coefficient of Figure 2, it is possible to estimate the overall effect of inventory on sales—which corresponds to the direct effect β_I plus the indirect effect through variety, $\beta_V\gamma$ (see Figure 2). In this case we are interested to estimate

$$Sales = (\beta_I + \beta_V\gamma)Inventory + \gamma_S' Z + \varepsilon'. \quad (3)$$

The error term ε' represents factors that affect sales that are unobservable in the data. Under the exogeneity assumption $E(\varepsilon' \mid Z, W)$, the coefficient $\beta_I' \equiv \beta_I + \beta_V\gamma$ can be estimated via instrumental variables, instrumenting *Inventory* with the weather variables W. This provides an alternative estimate of the overall effect of inventory on sales. The drawback of this approach is that it does not identify separately the effect of inventory and variety on sales. In particular, this precludes analyzing the allocation strategies described in Section 7.

Table 6. First-Stage Inventory Analysis—Weather Factors

	Inventory	*Variety*
Thunder	1.309***	0.008***
	(0.088)	(0.002)
Heat	0.346*	−0.007*
	(0.164)	(0.003)
Wind	−1.210***	0.018***
	(0.112)	(0.002)
Winter	−0.737***	−0.001
	(0.057)	(0.001)
Flood	1.679***	0.004*
	(0.101)	(0.002)
Tropical	0.910**	−0.114***
	(0.324)	(0.006)
*Variety D*1	3.974***	0.055***
	(0.121)	(0.002)
*Variety D*2	3.671***	0.049***
	(0.118)	(0.002)
*Variety D*3	3.037***	0.054***
	(0.111)	(0.002)
Dealer–model fixed effects	Yes	Yes
Dummies for each week and region	Yes	Yes
Dealer's local weather	Yes	Yes
Observations	216,438	216,438
Dealer–models	9,663	9,663

Note. Robust standard errors in parentheses.
$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

5. Results

To validate our instruments beyond the argument presented earlier, we present the results obtained on the first stage of the 2SLS estimation for each one of the two endogenous variables. This set of results include the six extreme weather factors described before and the variety measure for the three closest dealers, outside of the 150-mile exclusion area, from the focal dealer. Table 6 reports the estimation results for the first-stage analysis. The first column shows the estimates for *Inventory*, and the second column presents the estimation when instrumenting *Variety*. In addition to the relevant coefficients being significant in the first stage, we implemented an F-test of the instruments and verified that they are in a range that allows us to consider them statistically valid: 211.07 ($p = 0.000$) for the weather instrument in the *Inventory* estimation and 717.46 ($p = 0.000$) for the variety instruments in the *Variety* estimation. Both of these test results are well above the reasonable expectation for an instrument to be statistically meaningful.

Table 7 reports the results of the 2SLS estimation with bootstrap robust standard errors.[3] Column (1) shows the estimates of Equation (1), without instrumenting inventory or variety. Column (2) shows the estimates of Equation (1) after instrumenting the endogenous variables with the 2SLS approach. If the estimation does not account for potentially endogenous variables (column (1)), inventory seems to have a positive effect on sales, and the same is true for variety. However, the 2SLS estimates suggest that the direct effect of inventory (β_I in Figure (2)) is negative and statistically significant, but the effect of variety (β_V) is positive and also statistically significant. In other words, the 2SLS estimates indicate that sales increase if new submodels are made available to customers, but sales decrease if inventory is added to a submodel that is already available at the dealership.

Our estimates suggest that adding inventory without increasing variety has negative and statistically significant effect on sales. However, different vehicle allocation policies can give different results. Figure 3 illustrates the overall impact of inventory on sales with the vehicle allocation policy that maximizes the expansion of variety (solid line) compared to the allocation policy that expands inventory without increasing the number of submodels available (dashed line). As is apparent from the figure, whether adding inventory increases

Table 7. Main Model Results

	(1)	(2)	(3)	(4)	(5)
Inventory	0.0115***	−0.008***	−0.008***	0.013***	−0.009**
	(0.000)	(0.002)	(0.002)	(0.000)	(0.003)
Variety	0.308***	0.478***	0.621***		
	(0.011)	(0.118)	(0.047)		
Dealer–model fixed effects	Yes	Yes	Yes	Yes	Yes
Dummies for each week and region	Yes	Yes	Yes	Yes	Yes
Dealer's local weather	Yes	Yes	Yes	Yes	Yes
Observations	216,438	216,438	216,438	216,438	216,438
Dealer–models	9,663	9,663	9,663	9,663	9,663

Notes. (1) Estimation results for the impact of inventory and variety *without* instrumenting inventory or variety. (2) Estimation results for the impact of inventory and variety instrumenting the endogenous inventory and variety. (3) Estimation results for the impact of inventory and variety instrumenting *only* the endogenous inventory. (4) Estimation results obtained *without* instrumenting inventory. (5) Estimation results obtained instrumenting inventory. Robust standard errors in parentheses.
$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

Figure 3. (Color online) Sensitivity Analysis

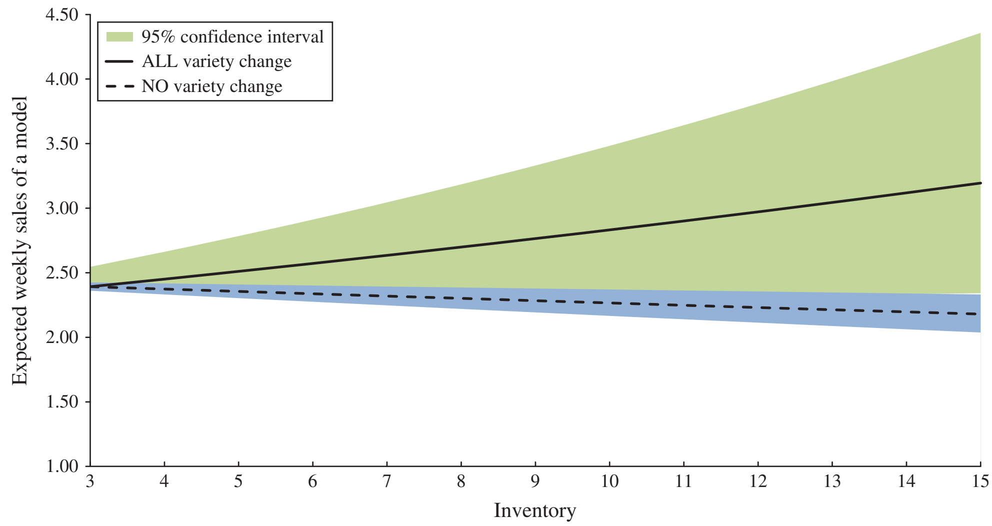

Note. This figure illustrates the overall impact of inventory on sales with the vehicle allocation policy that maximizes the expansion of variety (solid line) compared to the allocation policy that expands inventory without increasing the number of submodels available (dashed line), for a dealer that starts with three vehicles of a particular model.

or decreases overall sales depends on how vehicles are allocated to dealerships. A more precise analysis of alternative vehicle allocation policies is described in Section 7.

As proposed earlier, the overall impact of inventory on sales can be obtained by estimating Equation (3) directly. Column (4) of Table 7 shows these estimates when inventory is not instrumented, and column (5) shows the result when implementing the 2SLS approach for inventory alone. Once again we can observe how the instrument corrects the bias in the estimation when the endogeneity is ignored. The coefficient of inventory is −0.009, which is close to our previous estimate based on the coefficients of column (2) (which gave −0.008). This provides support to validate the consistency of the estimates of column (1).

As mentioned in Section 4, if inventory is set in anticipation of demand, then ε is likely to be positively correlated and therefore the OLS estimate of the inventory coefficient could be biased upward. Column (2) of Table 7 is consistent with this result: in fact, the bias is so severe that the coefficient on inventory changes sign and becomes negative with statistical significance. Column (3) of Table 7 shows the estimates of Equation (1) after instrumenting only the endogenous inventory and assuming that variety is not endogenous. We can observe that ignoring the potential bias of variety results in a larger estimate for variety, the difference between the two coefficients is statistically significant, suggesting that it is important to instrument both inventory and variety to obtain reliable estimates.

To repeat, the estimates in column (2) of Table 7 and the sensitivity analysis presented in Figure 3 suggest that (i) adding inventory decreases sales if variety is held constant (a scarcity effect); (ii) although increasing inventory can expand variety and variety has a positive impact on sales, the overall effect of increasing inventory is negative given the way vehicles are allocated in our sample; and (iii) adding inventory while simultaneously expanding variety can increase sales.

Mechanism. Several of the mechanism discussed earlier are consistent with these findings and several are not. For example, our findings are consistent with the notion that more variety improves the match between consumer preferences and the available inventory, thereby increasing the likelihood that a customer makes a purchase. In contrast, the results are not consistent with the notion that more variety creates confusion or higher evaluation costs, thereby reducing demand—in some categories it is possible that the confusion effect is real and sufficiently strong, but with automobiles it appears that consumers are more likely to buy when they have more options to choose from.

Our findings suggest that dealer pricing or consumer bargaining do not have a strong impact on the relationship between inventory and sales. As shown by Moreno and Terwiesch (2015), one would expect that a dealer is more likely to offer a better price when the dealer has an above-average amount of inventory because the dealer would want inventory to return to a more normal level. We observe that sales decrease as inventory increases (holding variety constant)—if this is to be explained by pricing, then one needs to be willing to assume that dealers increase their prices when they have more inventory. Similarly, our estimates cannot simply be explained by a stockout effect—if adding inventory prevents stockouts, then coefficient β_I should be positive, not negative.

Although our results indicate the presence of a scarcity effect, they are not consistent with all mechanisms that lead to a scarcity effect. For example, a scarcity effect can occur if consumers infer that ample inventory is a signal that a car is not popular, possibly because of poor design or quality. For this to explain our data, the inventory signal would have to be at the submodel level rather than at the model level—a consumer would have to believe that ample inventory of two-door Malibus is a bad signal for two-door Malibus, but the overall number of Malibus is not a negative signal. While we cannot rule this out, it does not seem plausible. We suspect that a consumer would infer quality, popularity, and design based on the total inventory of a model level rather than based on the inventory of each of the various submodels. If that is the case, then inferences of popularity cannot explain the negative relationship between sales and inventory, controlling for variety.

The scarcity effect we observe is consistent with the notion that inventory influences consumer search. Consumers are likely to desire a particular submodel. If there is only one unit available of their desired submodel, then they may discontinue their search for a new vehicle and purchase the vehicle. However, if the dealer has several units that fit the consumer's preference, the consumer may continue her search (or delay the purchase commitment to give further time to reflect on the decision), feeling confident that if she does not find a better match, she can return to the dealership. If the consumer continues her search, then at the very least it delays the sale, but worse, it risks losing the sale—the consumer might discover a better match at another dealership. Thus, we find evidence that low inventory reduces consumer procrastination and motivates an immediate sale.

6. Robustness Analysis

This section considers several robustness checks. For clarity in the exposition, we group these robustness analyses into two different categories: (1) alternative specifications to validate the identification strategy, and (2) accounting for possible interactions between dealerships and pricing.

6.1. Alternative Specifications

This section describes alternatives specifications that account for different seasonal controls, spurious effects, different specifications for the instruments (including additional plants located in Canada), and further validation on the variety instruments.

6.1.1. Seasonal Controls. Our identification is driven by variation in the weather events at the plants where vehicles were assembled and shipped from to the dealerships. Our estimation strategy would be invalidated if there is an unobserved factor that affects the weather at the plant and the likelihood of moving inventory produced at those plants to the dealers. However, extreme weather events are a natural experiment: there are no factors related to demand or supply of vehicles that could possibly correlate with the occurrence of these low-probability events at a specific plant. In other words, demand for vehicles does not plausibly cause weather events. That said, we must rule out correlation between potential demand confounders that are likely to occur at the same time of the year as weather events. Consequently, we validate our results considering multiple specifications that control for different levels of seasonality and trends.

In all of the following models, we consider the same structure as in the main model (a regression of sales as a function of two endogenous variables, *Inventory* and *Variety*—with the same instrumental variables) but we change the controls for trends and seasonality. We consider five additional specifications to the base model:

- Model A: Weekly dummies interacted with four different zones—our base model (104 dummies).
- Model B: Individual weekly dummies and a dummy for each one of the zones (30 dummies).
- Model C: Weekly dummies interacted with 26 different vehicle models (676 dummies).
- Model D: Individual weekly dummies and a dummy for each one of the modes (52 dummies).
- Model E: Individual linear trend for each one of the 26 vehicle models (26 separate trends by model).
- Model F: Weekly dummies and individual model trend (26 dummies and 26 separate trends).

Table 8 presents the result of the second stage under each one of these models. Column (1) corresponds to our base model (the one presented in column (2) of Table 7).

All six models present consistent results with respect to the magnitude and direction of our variables of interest. The results are robust to alternative specifications of seasonality, suggesting that the effect is not driven by unobservable seasonal effects, ruling this out as a potential confounder.

6.1.2. Spurious Effects. An additional concern is that, because the weather events studied are relatively infrequent, each event is influential in the sample and could potentially identify some spurious relationship. We implemented the following "placebo test" to rule out this explanation. We created 100 replications of our base model where we randomly generate the instruments for the analysis. The instruments were generated with the same frequency that appears in our data, and we force the event to occur within the same week range where we observe the events. We then estimated the model for each of the 100 replications, each one with a different set of randomly generated instruments. Out

Table 8. Alternative Seasonal Controls

	(1)	(2)	(3)	(4)	(5)	(6)
Inventory	−0.008***	−0.007**	−0.007**	−0.014***	−0.014**	−0.028***
	(0.002)	(0.002)	(0.002)	(0.004)	(0.004)	(0.005)
Variety	0.478**	0.466**	0.466**	1.207*	1.207*	1.610**
	(0.158)	(0.160)	(0.160)	(0.555)	(0.555)	(0.565)
Dealer–model fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Seasonal control	A	B	C	D	E	F
Dealer's local weather	Yes	Yes	Yes	Yes	Yes	Yes
Observations	216,438	216,438	216,438	216,438	216,438	216,438
Dealer–models	9,663	9,663	9,663	9,663	9,663	9,663

Notes. (A) Base model. Weekly dummies interacted with four different zones (104 dummies total). (B) Individual weekly dummies and a dummy for each one of the zones (30 dummies total). (C) Weekly dummies interacted with 26 different vehicle models (676 dummies total). (D) Individual weekly dummies and a dummy for each one of the modes (52 dummies total). (E) Individual trend for each one of the 26 vehicle models (26 separate trends by model). (F) Weekly dummies and individual model trend (26 dummies total and 26 separate trends). Robust standard errors in parentheses.

*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

of the 100 runs, in only seven instances are all instruments statistically significant at the 5% level, and the joint test of significance for the first stage of the instruments is larger than with the real instruments. Overall, this analysis passes the placebo test, providing further confidence that our instruments are valid and the main findings are not a result of pure chance.

6.1.3. Alternative Definition of Weather as an Instrument. We considered an alternative measure of the weather impact at the plants by using actual weather factors, rather than taking the definition of storm events from NOAA. To do this we constructed a set of extreme weather measures from the different weather factors captured by the different weather stations. These weather variables are described in detail in Table 2. We included *Wind*, *Fog*, *Rain*, and *Snow* variables because each of these weather events may influence travel to and from a plant. *Cloud* could proxy for other inclement weather and could influence employee behavior. *High Temp* is included because it could influence ambient temperature within the plant or employees that must work outside (e.g., loading docks). *Low Temp* may proxy for hazardous road conditions (e.g., ice). Some of the variables, such as *Wind* and *Cloud*, directly capture weather shocks. For other measures—specifically for *Rain*, *High Temp*, *Low Temp*, and *Snow*—we estimated specifications including multiple levels of the variable to capture potential nonlinear effects on production.

Some of these weather variables have a weak impact on dealership inventory, in part because of the high correlation between the many alternative measures of weather that we considered. Using a large number of instruments in a two-stage least square estimation can induce bias on the estimates (Buse 1992). There is also a rich literature that discusses other challenges that can arise when dealing with multiple instruments, in particular when some of these instruments might be weak (Bekker 1994, Donald and Newey 2001, Chao and Swanson 2005). Kloek and Mennes (1960) proposed a practical solution to solve the shortcomings of dealing with a large number of (possibly weak) instruments. The idea is to use a reduced number of principal components of the original set of instruments as the instrumental variables in the estimation. We follow a similar methodology in our approach.

We reduced the 13 weather variables to five principal components capturing more than 50% of the variance on the original variables. Hence, the components obtained contain a good portion of the information in our instruments.

The first stage of the 2SLS estimation of Equation (1) shows that when instrumenting *Inventory*, the five principal components coefficients are significant for the five factors—the average t value for each one of the factors is 12 and the $p-value < 0.000$ in every case. Once again, to validate the overall strength of our instrument, we observe that both the R-squared (0.9) of this regression and the F-test (184; $p < 0.000$) of join significance of the instruments exceed the usual standards to rule out weak instruments. To instrument *Variety*, we considered the same set of instruments included in our main results.

The results of the two 2SLS estimations with this set of instruments are included in Table 9. Column (1) shows the results when instrumenting both *Inventory* and *Variety*, which is analogous to the one presented in column (2) of Table 7. Column (2) of Table 9 shows the results when only *Inventory* is included in the analysis, as done for the analysis presented in column (5) of Table 7.

These results are directionally consistent with our main model results presented in Table 7, and the magnitude of the estimated effect is also comparable. This

Table 9. Alternative Weather Instruments and Additional Plants

	(1)	(2)	(3)	(4)
Inventory	−0.013**	−0.015***	−0.012**	−0.018***
	(0.004)	(0.003)	(0.004)	(0.003)
Variety	0.764**		0.764**	
	(0.304)		(0.269)	
Dealer–model fixed effects	Yes	Yes	Yes	Yes
Dummies for each week and region	Yes	Yes	Yes	Yes
Dealer's local weather	Yes	Yes	Yes	Yes
Observations	216,438	216,438	290,906	290,906
Dealer–models	9,663	9,663	12,853	12,853

Notes. (1) Estimation results for the impact of inventory and variety instrumenting inventory with weather factors. (2) Estimation results for the impact of inventory instrumenting inventory with weather factors. (3) Estimation results for the impact of inventory and variety instrumenting inventory and variety with weather factors. Includes additional plants. (4) Estimation results for the impact of inventory instrumenting inventory with weather factors. Includes additional plants. Robust standard errors in parentheses.

$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

gives us reassurance that the weather instruments capture an underlying mechanism and are not picking up a spurious effect. If anything, the results with this new set of instruments show a larger effect to the one presented with our main model since in this case the effect of inventory is more negative while the impact of variety is larger (comparing column (2) of Table 7 and column (1) of Table 9). However, because the criteria for constructing this new set of instruments are more subjective, we present our findings from the previous section as our main findings (which rely on the definition of extreme weather events provided by NOAA).

6.1.4. Further Validation of Variety Instrument. Our identification strategy to instrument variety is based on exploiting variation in the product variety of the shipments to the dealerships, on the premise that this variation arises from batching in the production and distribution of vehicles. If the plants were to operate under a stable production pattern, with smooth switching among the multiple submodels, the produced variety should be relatively stable over time and the shipments to dealerships should have a variety mix that is constant (hence weakening our proposed instrumental variable). To validate our approach, we conducted several data analyses to test if the shipments to the dealerships across weeks show a stable pattern of variety. We computed a variety measure as the fraction of submodels (of the total number of submodels of a model) that was shipped on each week (that is, with shipments greater than zero). If variety of a model is stable, this measure of variety should have a small variation across weeks. In contrast, the coefficient of variation for most models is between 0.15 and 0.30, which is considerable. Hence, the sample of shipments considered in the dealerships of this study exhibit a significant amount of variation in the product variety, suggesting a shipment process that is based on batches of submodels.

This analysis can be complemented with formal statistical tests. On each week, we computed the fraction of new shipments of a model corresponding to each of the submodels. Under the null hypothesis that the fraction of each submodel is stable, a confidence interval was computed to test if the observed fraction deviates significantly from the overall proportion of each submodel. These tests reveal that about 32% of the observations fall outside of the 99% confidence interval, rejecting the null that the fraction of shipments is stable. A similar test was conducted for each of the submodels, and the null of stable shipments is rejected for all of them.

To summarize, the data suggest that the actual flow of submodels to dealerships is not stationary across weeks, as the variation in product variety shipped to the dealerships cannot be explained by pure random chance. The batching pattern of shipments to dealerships can be used as an exogenous source of variation in product variety that is used to identify the effect of this variable.

6.1.5. Additional Plants. Constructing the weather instruments, as described in the previous section, allows us to include additional production plants in our analysis. GM has a number of plants in Canada that produce car models that are sold in the United States. NOAA does not report extreme weather events, our preferred instrumental variable, for Canada. However, the Canadian weather stations offer the historical information for the same set of weather variables we considered to estimate the weather factors in the previous section. This allows us to expand our sample of both plants and dealer–models and further validate our estimation results.

We estimated a new set of five weather factors implementing a similar approach to the one described in the previous section. Using these factors, we instrumented inventory and sales and included the additional models. The results of this analysis are presented in Table 9. Column (3) shows the results when both *Inventory* and *Variety* are included, and column (4) presents the results when we consider only the impact of *Inventory* on sales. Again, these results are directionally consistent and of the same magnitude as our main model results presented in Table 7 (columns (2) and (5)). This indicates that our findings are robust to an expanded sample with additional plants and models.

6.2. Dealers' Interactions

We consider competition among dealers, the proportion of dealers' inventory across different model years, sales driven by transfers, and transfers between dealers.

6.2.1. Competition Among Dealers. The dealers in our sample face different levels of competition from GM and non-GM dealerships. As mentioned earlier, the dealer–model fixed effects included in our main specification account for the average competition intensity for a particular model at a dealer. However, the inventory level for a model could vary across dealers from one week to another, and this variation may be known to the dealers. To explain our results, low inventory at dealer A in a market would need to be correlated with reduced competition from the other dealerships in the same market, thereby allowing dealer A to increase his sales. Although we do not view this as likely, to explore the impact of competition, we estimate our main model with a subsample of dealers that do not face competition in their local market from other GM dealers. Based on empirical work defining the relevant market for a dealership (Albuquerque and Bronnenberg 2012), we defined a subsample of dealers with no competing GM dealer (of any GM brand) within a 15-mile radius.

Table 10 reports the analysis with this subsample. Columns (1) and (3) correspond to the analysis ignoring the bias, and columns (2) and (4) correspond to the analysis when we implement the 2SLS with the extreme weather instruments. These results are consistent with the results obtained with the complete sample and suggest that our main results are not confounded by the impact of competition patterns between GM dealers.

6.2.2. Inventory Age. Our analysis does not have controls to capture pricing behavior. It is plausible that price markdowns occur more frequently for older models, and so part of the markdown pricing can be captured by including controls that measure the age of the inventory. In our sample, only 4.29% of vehicles are from the 2006 model year and the rest are 2007 model-year vehicles, suggesting that the effect of inventory age may be small. Nevertheless, we report in column (4) of Table 11 the estimation results from our main specifications with the additional control of the percentage of 2006 model-year vehicles in inventory for each model, dealer, and week. There is a positive but statistically insignificant coefficient on the variable capturing the age of inventory. More importantly, the main effects of interest (inventory and variety) are similar and statistically identical as in the main results. In conclusion, the age of inventory has no effect on the estimates, perhaps because of the low prevalence of older inventory in our sample.

Table 10. Isolated Dealers Analysis

	(1)	(2)	(3)	(4)
Inventory	0.010***	−0.016***	0.012***	−0.020**
	(0.000)	(0.004)	(0.000)	(0.006)
Variety	0.295**	0.784**		
	(0.015)	(0.290)		
Dealer–model fixed effects	Yes	Yes	Yes	Yes
Dummies for each week and region	Yes	Yes	Yes	Yes
Dealer's local weather	Yes	Yes	Yes	Yes
Observations	112,177	112,177	112,177	112,177
Dealer–models	5,131	5,131	5,131	5,131

Notes. (1) Estimation results for the impact of inventory and variety *without* instrumenting inventory or variety. (2) Estimation results for the impact of inventory and variety instrumenting the endogenous inventory and variety. (3) Estimation results obtained *without* instrumenting inventory. (4) Estimation results obtained instrumenting inventory. Robust standard errors in parentheses.

$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

Table 11. Transfers and Model-Year Robustness

	(1)	(2)	(3)	(4)
Inventory	−0.008***	−0.008***	−0.008***	−0.008***
	(0.002)	(0.002)	(0.002)	(0.002)
Variety	0.478**	0.468**	0.454**	0.491**
	(0.158)	(0.166)	(0.157)	(0.159)
TransSale			0.231**	
			(0.004)	
Model06				0.007
				(0.007)
Dealer–model fixed effects	Yes	Yes	Yes	Yes
Dummies for each week and region	Yes	Yes	Yes	Yes
Dealer's local weather	Yes	Yes	Yes	Yes
Observations	216,438	216,438	216,438	216,438
Dealer–models	9,663	9,663	9,663	9,663

Notes. (1) Base case. Dealers outside of a 150-mile radius for the variety instrument. (2) Dealers outside of a 200-mile radius for the variety instrument. (3) Includes dummy variable (*TransSale*) corresponding to sales from transfers. (4) Includes variable *Model06*, the percentage of vehicles corresponding to 2006 models. Robust standard errors in parentheses.

$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

6.2.3. Sales Driven by Transfers. We evaluate our main specification, introduced in Equation (1), where we include a dummy variable to explicitly control whether a sale, on a particular dealer–model, was from a transferred vehicle. The results of this analysis are presented in column (3) of Table 11. As expected, the indicator variable noting whether a sale was from a transfer is statistically significant because it is correlated with a sale by definition. More importantly to the focus of our study, the results for our variables of interest hold under this alternative specification.

6.2.4. Transfers Between Dealers. If a dealer lacks a submodel that a consumer wants, the dealer can try to convince the consumer to purchase a different submodel or the dealer can try to find the desired submodel at a nearby dealer. If the desired vehicle is found at another dealer, a transfer can occur between the two dealerships if they can agree to the transfer. In many cases this transfer involves a swap of vehicles rather than an exchange of cash—the requesting dealer has a customer that will purchase the donating dealer's vehicle, and in return, the requesting dealer offers one of its vehicles to the donating dealer, who probably does not have a customer at the time of the swap willing to purchase the vehicle received. In our sample, 12.9% of total sales are from vehicles that were transferred from one dealership to another. We next discuss three issues raised by the presence of transfers.

The use of transfers could be consistent with a negative relationship between inventory and sales—when a dealer has fewer vehicles, it compensates by making more frequent use of transfers. Through this mechanism, lower inventory does not induce higher sales because of influencing consumer search or preferences, but rather by influencing dealer behavior (i.e., they do more transfers). To test this alternative hypothesis, we considered a model with the log of transfer sales as the dependent variable and inventory at the dealer as the independent variable, and the same set of controls included in Z. This regression of transfer sales does not find a significant relationship between inventory and transfer sales, which suggests our results are not well explained by the additional use of transfers when inventory is low.

We do not observe all transfers because while we know the location of all GM dealerships within the country, we could not monitor the inventory of dealerships in all states in the United States. Hence, there are some unobserved transfers between dealers across states. To address this concern, we identify for each state a radius such that 90% of the transfers in our sample occur between dealerships within this radius. Across the entire sample, this radius is 200 miles (i.e., 90% of transfers in our sample occur between dealers that are within 200 miles of each other). As expected, the radius is higher for less dense states such as Nebraska and lower for states like California. Next, we restrict our sample to dealerships for which at least 75% of the dealership's neighbors within the state-specific radius are included in our set of monitored dealerships. For example, a dealership in northern Florida might be excluded because only 60% of the dealerships within 196 miles (the radius for Florida) are within Florida. This excludes dealerships that are more likely to have unobserved transfers (because they are close to state borders) and retains dealerships for which we are likely to observe nearly all of their actual transfers. The results with the restricted sample do not change in sign or magnitude, suggesting that unobserved transfers do not affect our results.

Finally, the use of transfers is relevant for the validity of the variety instrument, which requires that the variety of distant dealers is correlated with a focal dealer's variety but not the focal dealer's sales through unobserved variables. This would be a concern if a focal dealer used the availability of distant vehicles to increase its sales. The fact that inventory is uncorrelated with the use of transfer encourages us that this is not the case. In addition, dealers generally attempt to sell what is on their lot, suggesting that they do not encourage customers to select from distant inventory. But for further confidence, we replicated the main model using a variety instrument based on a 200-mile exclusion instead of the main model's 150-mile radius. Only 10% of transfers come from dealers that are 200 miles or greater away, and only 1.8% of those are from dealers used in the variety instrument. Hence, the vehicles used for our variety instrument represent only 0.166% of all vehicles in our sample. Table 11 reports that our results are essentially unchanged using this more restrictive sample for the variety instrument. This gives us confidence that our variety instrument is valid even in the presence of transfers.

7. The Impact of Inventory Allocation

Our empirical estimation reveals that adding inventory to a dealer is only beneficial if the added vehicle expands the dealer's set of submodels—increasing the inventory of a particular submodel actually lowers sales. This section explores the potential sales benefit of using this result to better allocate vehicles to dealers. We take two different approaches. The first approach estimates the potential sales improvement from reallocating the existing vehicles among the dealers in a small local area. The second approach considers only the incoming vehicles to a larger region (e.g., a state) and attempts to maximize sales by allocating those vehicles to the dealers in the area while leaving the dealers' existing inventory intact.

Given the size of our data set (1,289 dealers, 30 weeks, etc.), we focus our analysis on a particular week (the week with the median number of total cars) and the 10 most popular models. These models—Cobalt, Equinox, G6, HHR, Impala, Suburban, Tahoe, TrailBlazer, Saturn, VUE, and Yukon—represent approximately 60% of the sales across all of the GM models in our sample. We emphasize that this analysis is only suggestive of the potential benefit of changing their allocation process. Because of computational effort, we are unable to analyze all reasonable parameters estimates, all models, and all weeks in our sample.

7.1. Local Reallocation Among Dealers

The analysis in this section partitions dealers into small local markets. Dealership inventory can be observed in detail at the submodel level: some dealers may have multiple units within a submodel and other dealers within the same local market might not have any vehicles of that submodel. Hence, based on our results, both dealers could benefit from a vehicle transfer—moving a vehicle from the dealer with multiple units to the dealer with no units increases sales at both dealers. Thus, we evaluate for each model the total sales gain across all markets that could be achieved by efficient vehicle transfers so as to maximize the variety each dealer offers and to minimize the duplication of units within submodels. We do not model the cost of actually transferring these vehicles—any sales improvement from reallocation would have to be compared with the cost of achieving the better balance of variety across dealers.

We group dealers as part of the same local market if they are in the same core-based statistical area (CBSA).[4] We consider vehicle swaps only between dealers in the same CBSA, so that the total inventory within each CBSA remains constant. In addition, we impose a cap on the amount of inventory a dealership can reduce of a given model.

The decision variable is the number of vehicles of submodel k of model j allocated to dealership i, Q_{ijk}, in a selected week. For each CBSA and each model j, we solve the following integer nonlinear optimization problem:

$$\max_{Q_{ijk}} \left[\sum_{i=1}^{n} \exp\left(\hat{\delta}_{ij} + \hat{\beta}_I \sum_{k=1}^{m_j} Q_{ijk} + \hat{\beta}_V \cdot \text{Variety}_{ij} \right) \right] \tag{4}$$

s.t.

$$\sum_{i=1}^{n} Q_{ijk} = \sum_{i=1}^{n} I_{ijk} \quad \forall j, k, \tag{5}$$

$$\sum_{k=1}^{m_j} Q_{ijk} = (1-\lambda) \sum_{k=1}^{m_j} I_{ijk} \quad \forall i, j, \tag{6}$$

$$\text{Variety}_{ij} = \frac{\sum_{k=1}^{m_j} \mathbb{I}(Q_{ijk} \geq 1)}{m_j}, \tag{7}$$

$$0 \leq Q_{ijk} \leq T_{ijk}, \tag{8}$$

where n is the number of dealers within a CBSA; m_j is the total number of submodels for model j; I_{ijk} is dealer i's initial endowment of inventory of submodel jk (i.e., if there is no reallocation); $\hat{\beta}_V$ and $\hat{\beta}_I$ are the estimated coefficients from column (1) of Table 7; $\hat{\delta}_{ij}$ is the estimated fixed effect for dealer i and model j; and T_{ijk} is the maximum number of vehicles that dealer i carried for submodel jk.

Constraint (5) ensures that the reallocation does not change the total inventory within the CBSA of model j. Constraint (6) ensure that dealer i's inventory of model j cannot decrease more than a λ fraction of the initial inventory. When $\lambda = 0$, these constraints imply that each dealer that gains a vehicle must also give up a vehicle of the same model. Constraint (8) ensures that at the end of the swaps the maximum number of units of a particular submodel k at dealer i will be less or equal to the maximum number of units of any submodel k that dealer i was carrying at the beginning of the swaps. The objective is then to maximize $Variety_{ij}$ while keeping each dealership's inventory constant.

Table 12. The Impact of Inventory Allocation

	CBSA reallocation			State reallocation
	No inventory reduction allowed	10% inventory reduction allowed		
	Sales increase (%)	Inventory reduction (%)	Sales increase (%)	Sales increase (%)
Cobalt	4.2	1.2	4.5	2.9
Equinox	4.6	2.4	5.0	2.2
G6	5.0	1.8	5.3	2.3
HHR	5.1	4.7	5.8	6.8
Impala	5.9	2.0	6.4	5.1
Suburban	4.6	1.4	4.9	1.8
Tahoe	3.9	2.5	4.6	1.9
TrailBlazer	4.8	1.2	5.3	1.0
Saturn VUE	1.1	9.1	5.0	2.2
Yukon	3.9	3.8	5.1	2.5
Weighted average	4.4	2.8	5.2	2.8

The first column of Table 12 shows the solution to this math program when imposing the constraint that inventory of all model–dealers remains constant ($\lambda = 0$), as measured by the average potential sales improvement. We find that on average, exchanging inventory among dealers within a CBSA with the objective of maximizing each dealer's offered variety yields a weighted average sales gain of 4.4%.

The second and third columns of Table 12 allow up to a 10% reduction of a model's inventory at each dealer ($\lambda = 0.1$). The second column shows the average inventory reduction that takes place at the dealers for each model. The third column presents the results for the average potential sales improvement for each car model. We find that exchanging inventory among dealers within a CBSA while giving the dealers the option to reduce their inventory results in a reduction on average of 2.8% of the dealer's inventory and weighted average sales gain of 5.2%.

7.2. Statewide Reallocation of Vehicles

Instead of swapping vehicles after they arrive at dealerships, we now consider changing the allocation of vehicles after they leave the production facility. At that point in time there may be some flexibility with respect to a vehicle's final destination, and this flexibility may come with little incremental cost. In particular, we estimate the sales gain that can be achieved through smarter allocation of vehicles that arrive to a particular state in a given week. With this approach there are no transfers among dealers—each dealer's initial inventory remains with that dealer. However, rather than send submodel k to a dealer who already has some units of submodel k, it is better to send that vehicle to a dealer who begins the week without any units of submodel k. Let $Y_{ijk} > 0$ denote the number of incoming vehicles of model j/submodel k that are allocated to dealer i. The resulting math program is similar to (4)–(8) but replacing constraints (5)–(6) by

$$\sum_{i}^{n} Y_{ijk} = A_{jk}, \tag{9}$$

$$\sum_{k}^{m_j} Q_{ijk} \leq M_{ij}, \tag{10}$$

$$Q_{ijk} = I_{ijk} + Y_{ijk}. \tag{11}$$

Constraint (9) ensures that the state receives the same number of vehicles of model j and submodel k as we observed in our data for the chosen week. Constraint (10) ensures that dealer i's inventory of model j after the assignment is not greater than the maximum number of vehicles of model j that dealer i had in any week of our sample. This precludes allocations that result in some dealers having an unreasonably large amount of inventory. Equation (11) states that the resulting dealer's inventory of a model equals the dealer's initial endowment, I_{ijk}, plus the dealer's allocation, Y_{ijk}.

The third column of Table 12 shows the results for each model in this statewide allocation problem. On average, we find that routing vehicles to dealers in a state so as to minimize overlap within a dealer's inventory while maximizing variety across dealers yields an average sales increase of 2.8%.

8. Conclusion

We developed an econometric model to estimate the effect of inventory and variety on sales at U.S. automobile dealerships. There exist contradictory theories with respect to the impact of inventory on sales. There are several mechanisms that lead to a billboard effect—a positive relationship between inventory and sales. For example, at a basic level, adding inventory can increase sales by reducing stockouts, or by expanding the variety of submodels available. However, there are mechanisms that lead to a scarcity effect—a negative relationship between inventory and sales. For instance, adding inventory may encourage additional search. In our sample, given how vehicles were allocated to dealerships, we find that an increase in inventory at a dealer actually lowers sales. Furthermore, we decompose this effect into two parts: (1) increasing inventory of a submodel does indeed reduce sales, but (2) if increasing inventory expands the number of submodels available to customers, then sales increase. In short, the benefit of expanding variety can dominate the negative effect of increasing inventory within a submodel. This result emphasizes that the total level of inventory can be a poor proxy for the presence of either a billboard or scarcity effect.

Our finding is consistent with two mechanisms relating inventory to sales: (i) expanded variety enables a better fit to consumer preferences, thereby increasing sales; and (ii) too many of the same submodel encourages consumers to procrastinate in their purchase decision, thereby lowering sales. To maximize sales, a dealer wants to have one unit of each submodel (to generate an urgency to "buy now before they are all gone") while also having as many submodels available as possible, to cater to the heterogeneous tastes of consumers.

Our results indicates that it is important to consider how vehicles are allocated. The data suggest that vehicles are allocated in a way that does not maximize the heterogeneity of submodels available to consumers. Dealers may view one submodel as particularly desirable and then take actions to increase their inventory in that submodel rather than to expand the set of submodels offered. For example, based on our estimates, an allocation policy that is focused on maximizing variety can increase sales by about 4.4%, without changing the number of vehicles produced or the number of vehicles each dealer carries. In other words, through a "maximize variety, minimize duplication" allocation strategy, it is possible to increase sales with no additional cost.

Endnotes

[1] We also conducted the analysis restricting the sample by dropping dealers with more than 75% unobserved neighbor dealers, and the results were similar.

[2] The locations consider for this analysis were Marysville and Columbus, Ohio; Washington DC and Baltimore, Maryland; Kansas City, Missouri, and Topeka, Kansas; and Lansing and Grand Rapids, Michigan.

[3] The bootstrap was implemented at the dealer–model level. We select a random sample of dealer–model combinations and include the complete panel for the selected pairs. We implemented 500 bootstraps for our main results.

[4] CBSA is a U.S. geographic area defined as an urban center of at least 10,000 people and adjacent areas that are socioeconomically tied to the urban center by commuting.

References

Albuquerque P, Bronnenberg BJ (2012) Measuring the impact of negative demand shocks on car dealer networks. *Marketing Sci.* 31(1):4–23.

Aviv Y, Pazgal A (2008) Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing Service Oper. Management* 10(3):339–359.

Balachander S, Liu Y, Stock A (2009) An empirical analysis of scarcity strategies in the automobile industry. *Management Sci.* 55(10):1623–1637.

Balakrishnan A, Pangburn MS, Stavrulaki E (2004) "Stack them high, let 'em fly": Lot-sizing policies when inventories stimulate demand. *Management Sci.* 50(5):630–644.

Barkholz D (2010) Dealers: We want more cars! In quest for profits, factories keep inventory tight. *Automotive News* (August 23), http://www.autonews.com/apps/pbcs.dll/article?AID=/20100823/OEM01/308239949/ 1262.

Bekker PA (1994) Alternative approximations to the distributions of instrumental variable estimators. *Econometrica* 62(3):657–681.

Bernstein F, Federgruen A (2004) A general equilibrium model for industries with price and service competition. *Oper. Res.* 52(6):868–886.

Brehm SS, Brehm JW (1981) *Psychological Reactance: A Theory of Freedom and Control* (Academic Press, New York).

Brock TC (1968) *Implications of Commodity Theory for Value Change* (Academic Press, New York).

Buse A (1992) The bias of instrumental variable estimators. *Econometrica* 60(1):173–180.

Busse MR, Pope DG, Pope JC, Silva-Risso J (2014) The psychological effect of weather on car purchases. *Quart. J. Econom.* 130(1):371–414.

Cachon GP, Swinney R (2009) Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Sci.* 55(3):497–511.

Cachon GP, Terwiesch C, Xu Y (2008) On the effects of consumer search and firm entry in a multiproduct competitive market. *Marketing Sci.* 27(3):461–473.

Chao JC, Swanson NR (2005) Consistent estimation with a large number of weak instruments. *Econometrica* 73(5):1673–1692.

Dana JD Jr, Petruzzi NC (2001) Note: The newsvendor model with endogenous demand. *Management Sci.* 47(11):1488–1497.

DeGraba P (1995) Buying frenzies and seller-induced excess demand. *RAND J. Econom.* 26(2):331–342.

Deneckere R, Peck J (1995) Competition over price and service rate when demand is stochastic: A strategic analysis. *RAND J. Econom.* 26(1):148–162.

Donald SG, Newey WK (2001) Choosing the number of instruments. *Econometrica* 69(5):1161–1191.

Gourville JT, Soman D (2005) Overchoice and assortment type: When and why variety backfires. *Marketing Sci.* 24(3):382–395.

Hoch SJ, Bradlow ET, Wansink B (1999) The variety of an assortment. *Marketing Sci.* 18(4):527–546.

Iyengar SS, Lepper MR (2000) When choice is demotivating: Can one desire too much of a good thing? *J. Personality Soc. Psych.* 79(6):995–1006.

Kalyanam K, Borle S, Boatwright P (2007) Deconstructing each item's category contribution. *Marketing Sci.* 26(3):327–341.

Kloek T, Mennes LBM (1960) Simultaneous equations estimation based on principal components of predetermined variables. *Econometrica* 28(1):45–61.

Kuksov D, Villas-Boas JM (2010) When more alternatives lead to less choice. *Marketing Sci.* 29(3):507–524.

Matsa DA (2011) Competition and product quality in the supermarket industry. *Quart. J. Econom.* 126(3):1539–1591.

Moreno A, Terwiesch C (2015) Pricing and production flexibility: An empirical analysis of the U.S. automotive industry. *Manufacturing Service Oper. Management* 17(4):428–444.

Murray KB, Muro FD, Finn A, Leszczyc PP (2010) The effect of weather on consumer spending. *J. Retailing Consumer Services* 17(6):512–520.

National Oceanic and Atmospheric Administration (2015) Storm data publication. Last accessed June 29, 2018, http://www.ncdc.noaa.gov/stormevents/ftp.jsp.

National Oceanic and Atmospheric Administration (2016) National Weather Service instruction. Last accessed June 29, 2018, http://www.nws.noaa.gov/directives/sym/pd01016005curr.pdf.

Qian L, van Ryzin GJ (2008) Strategic capacity rationing to induce early purchases. *Management Sci.* 54(6):1115–1131.

Schwartz B (2004) *The Paradox of Choice: Why More Is Less* (Ecco/HarperCollins Publishers, New York).

Smith SA, Achabal DD (1998) Clearance pricing and inventory policies for retail chains. *Management Sci.* 44(3):285–300.

Soysal GP, Krishnamurthi L (2012) Demand dynamics in the seasonal goods industry: An empirical analysis. *Marketing Sci.* 31(2):293–316.

Steele AT (1951) Weather's effect on the sales of a department store. *J. Marketing* 15(4):436–443.

Stock A, Balachander S (2005) The making of a "hot product": A signaling explanation of marketers' scarcity strategy. *Management Sci.* 51(8):1181–1192.

Su X, Zhang F (2008) Strategic customer behavior, commitment, and supply chain performance. *Management Sci.* 54(10):1759–1773.

Su X, Zhang F (2009) On the value of commitment and availability guarantees when selling to strategic consumers. *Management Sci.* 55(5):713–726.

Talluri K, van Ryzin G (2004) Revenue management under a general discrete choice model of consumer behavior. *Management Sci.* 50(1):15–33.

Tereyağoğlu N, Veeraraghavan S (2012) Selling to conspicuous consumers: Pricing, production, and sourcing decisions. *Management Sci.* 58(12):2168–2189.

Train KE (2009) *Discrete Choice Methods with Simulation*, 2nd ed. (Cambridge University Press, New York).

Worchel S, Lee J, Adewole A (1975) Effects of supply and demand on ratings of object value. *J. Personality Soc. Psych.* 32(5):906–914.

Zettelmeyer F, Morton FS, Silva-Risso J (2006) Scarcity rents in car retailing: Evidence from inventory fluctuations at dealerships. NBER Working Paper 12177, National Bureau of Economic Research, Cambridge, MA.