

How Market Power Affects Dynamic Pricing: Evidence from Inventory Fluctuations at Car Dealerships

Ayelet Israeli,^a Fiona Scott-Morton,^{b,c} Jorge Silva-Risso,^d Florian Zettelmeyer^{c,e}

^aHarvard University, Boston, Massachusetts 02163; ^bYale University, New Haven, Connecticut 06520; ^cNational Bureau of Economic Research, Cambridge, Massachusetts 02138; ^dUniversity of California, Riverside, Riverside, California 92521; ^eNorthwestern University, Evanston, Illinois 60208

Contact: aisraeli@hbs.edu,  <https://orcid.org/0000-0002-5996-3202> (AI); fiona.scottmorton@yale.edu (FS-M); jorge.silva-risso@ucr.edu (JS-R); f-zettelmeyer@kellogg.northwestern.edu,  <https://orcid.org/0000-0003-0791-2849> (FZ)

Received: April 19, 2018

Revised: August 14, 2020; December 8, 2020

Accepted: December 13, 2020

Published Online in Articles in Advance:
October 14, 2021

<https://doi.org/10.1287/mnsc.2021.3967>

Copyright: © 2021 INFORMS

Abstract. This paper investigates empirically the effect of market power on dynamic pricing in the presence of inventories. Our setting is the auto retail industry; we analyze how automotive dealerships adjust prices to inventory levels under varying degrees of market power. We first establish that inventory fluctuations create scarcity rents for cars that are in short supply. We then show that dealers' ability to adjust prices in response to inventory depends on their market power, that is, the quantity of substitute inventory in their selling area. Specifically, we show that the slope of the price–inventory relationship (higher inventory lowers prices) is significantly steeper when dealers find themselves in a situation of high rather than low market power. A dealership with high market power moving from a situation of inventory shortage to a median inventory level lowers transaction prices by about 0.57% *ceteris paribus*, corresponding to 32.5% of dealers' average per-vehicle profit margin or \$145.6 on the average car. Conversely, when competition is more intense, moving from inventory shortage to a median inventory level lowers transaction prices by about 0.35% *ceteris paribus*, corresponding to 20.2% of dealers' average per-vehicle profit margin or \$90.9. To our knowledge, we are the first to empirically show that market power affects firms' ability to dynamically price.

History: Accepted by Juanjuan Zhang, marketing.

Supplemental Material: Data and the online appendix are available at <https://doi.org/10.1287/mnsc.2021.3967>.

Keywords: marketing • pricing • price discrimination • inventory production • dynamic pricing • automobile industry • market power

1. Introduction

Since its initial success in airlines, dynamic pricing has become ubiquitous in many competitive industries, for example, cruise lines, apparel, rental car companies, and hotels. In conjunction with these applications, a rich academic literature has proposed models of dynamic pricing when firms have limited market power. Empirically, in contrast, we know little about how market power affects firms' ability to dynamically price. In this paper, we try to provide such empirical evidence.

Our setting is the auto retail industry. Dealers engage in dynamic pricing because car supply—in the short term—is restricted to the inventory on a dealer's lot, and demand is volatile. As a result, the opportunity cost of selling a car of a specific make, model, options, and color is continually changing with demand for that particular car within a geographic market.¹ Thus, there are effectively new, dealer-level optimal prices each day—or perhaps more frequently—for each vehicle. Negotiating with the customer allows the dealer to incorporate the latest information on inventory levels into the offered price.

There are two (related) sources of market power in auto retailing. First, a dealer's market power depends on the number of competing dealers within the selling area. Second, holding constant the number of competing dealers, a dealer's market power also varies with the quantity of substitute inventory available for sale by competing dealers. The number of competing dealers is stable in the medium run. In contrast, the amount of substitute inventory is quite volatile because it is subject to demand shocks.

In this paper, we empirically show that a dealer's ability to adjust prices in response to inventory depends on the second source of market power, that is, the quantity of substitute inventory in the selling area. We first show that inventories systematically affect pricing in the car retailing industry. Second, we show that the slope of the price–inventory relationship (higher inventory lowers prices) is significantly steeper when dealers find themselves in a situation of high rather than low market power.

We are not the first to point out empirically that competition or market power affects prices and inventories

when firms dynamically price. Amihud and Mendelson (1989) use public data to document that firms lower their inventories as their market power decreases (as measured by the firms' market shares and margins). Using automobile data, Cachon and Olivares (2010) show that, at the level of automotive brands, there is a positive relationship between the number of dealerships and inventory (among other findings). Using data in individual transactions at GM dealerships, Olivares and Cachon (2009) distinguish between sales and service effects of inventory. They show that the service effect leads dealers to carry more inventory (holding sales constant) when they face additional competition. Borenstein and Rose (1994) show that, in the airline industry, the more competitive a particular route is, the greater the price dispersion because of price discrimination.

Although these papers provide convincing evidence on the way that market power affects inventories or price, to our knowledge, we are the first to show empirically how market power changes the *price–inventory relationship*. Specifically, we show that firms' ability to adjust prices in response to inventory varies with market power. The paper closest ours is Siegart and Ulbricht (2020), who document that airline ticket fares increase over time prior to departure and that this increase is flatter in more competitive routes. However, the results of that paper are correlational because the competitiveness of routes is endogenous. In contrast, we are able to identify the effect of market power on the price–inventory relationship using exogenous intertemporal changes in substitute inventory.²

To illustrate why market power might affect the price–inventory relationship, it is helpful to understand why there is a price–inventory relationship in the first place. Consider a monopolistic dealer who can periodically reorder inventory but who faces a time delay between ordering and the arrival of inventory.³ If a dealer's inventory of a particular car increases, given some resupply schedule, the dealer's opportunity cost from selling that vehicle has decreased because the car is now less scarce relative to expected future demand. In contrast, when inventory is low, any sale has a higher opportunity cost because the dealer may not be able to sell to a future high-valuation customer who could arrive after the last car is sold but before the new inventory arrives. Notice that this reasoning holds even if the dealer is correct about the distribution from which the reservation prices of buyers are drawn; the argument does not depend on a dealer updating expectations or "learning" about the underlying level of demand.

To understand how the quantity of substitute inventory might affect this price–inventory relationship, consider the extreme case in which there is *no* substitute inventory—the situation a monopolist dealer faces. As described, lower inventory should lead to higher prices. Now, consider the other extreme case in

which there is ample *perfectly substitutable* inventory. The dealer is not able to raise prices when the dealer's inventory is very low; consumers can easily find substitute inventory at another dealer, making them much more price elastic.

In practice, the large number of options with which dealers order cars of the same make and model imply that "substitute inventory" at other dealers rarely represents a perfect substitute. Instead, consumers are more likely to find a close substitute to a focal vehicle the more substitute inventory other dealerships have in stock. As a result, we expect that the slope of the price–inventory relationship is smaller in magnitude the more inventory competing dealers have of the same make and model.⁴ In summary, we hypothesize that dealers have an incentive to engage in dynamic pricing, but their ability to do so is weakened as they face more competition.

The empirical section of the paper provides evidence for the hypothesized relationship between our market power measure—the quantity of substitute inventory—and the strength of the price–inventory relationship. We classify vehicle sales into quartiles, depending on the amount of substitute inventory that was available in the dealer's market area at the time of purchase. Not surprisingly, higher levels of substitute inventory are associated with lower average prices, and prices increase with market power. However, the level of substitute inventory also changes the price–inventory relationship at dealers. When there is a shortage of substitute inventory (quartile 1), a dealership moving from a situation of (own) inventory shortage to a median (own) inventory level lowers transaction prices by about 0.57% *ceteris paribus*, corresponding to 32.5% of dealers' average per-vehicle profit margin or \$145.6 on the average car. Conversely, when there is ample substitute inventory (quartile 4), moving from inventory shortage to a median inventory level lowers transaction prices by about 0.35% *ceteris paribus*, corresponding to \$90.9 or 20.2% of dealers' average per-vehicle profit margin. For quartiles 2 and 3, we find intermediate effects, at 0.51% and 0.43%, respectively. Overall, as hypothesized, dynamic pricing is more pronounced when dealers have more market power.

We consider the potential endogeneity of prices and inventory levels resulting from, for example, a temporary demand shock that raises the price of a model and lowers inventory levels. We use a series of fixed effects specifications as well as instrumental variables to control for this potential problem. Our results remain robust to these approaches as well as to alternative definitions of inventory and substitute inventory.

We also find that the price–inventory relationship extends to the margin a retailer obtains from financing and insurance (F&I margins). In particular, for below-median inventory levels (14 and fewer cars), one

additional car in inventory is associated with a margin that is lower by 0.005%. That is, a dealership moving from a situation of (own) inventory shortage to a median (own) inventory level lowers F&I margins by about 0.065%. For above-median inventory levels, the coefficient is only 0.0003% for each additional car. This F&I margin–inventory relationship also depends on market power as the slope is steeper when dealers have more market power. Therefore, dealers' ability to dynamically price financing and insurance options is weakened as the quantity of substitute inventory increases.

In addition to the empirical literature on the effect of market power on prices and inventory, there is a substantial theoretical literature on dynamic pricing in operations research and economics. In operations research, several papers show that competition changes prices in the presence of inventory (Dudey 1992, Xu and Hopp 2006, Anderson and Schneider 2007, Mookherjee and Friesz 2008, Mantin et al. 2011, Martínez-de Albéniz and Talluri 2011). Lin and Sibdari (2009) show that, under competition, the optimal price for a product need not be nondecreasing in time to go. Xu and Hopp (2006) show that firms may overstock in competitive situations relative to monopoly. Gallego and Hu (2014) formulate an intertemporal pricing problem under competition as a differential game and show that this formulation sheds light on how market conditions and supply constraints affect intertemporal pricing. Liu and Zhang (2013) and Levin et al. (2009) model dynamic pricing under competition when consumers are strategic. In economics, some research analyzes the interplay of prices and inventory, albeit with a substantially different focus than that of our paper. Hall and Rust (2000) analyze the pricing and inventory behavior of a steel wholesaler who also negotiates prices with customers and displays substantial fluctuation in day-to-day inventory of different products. Copeland et al. (2011) model the optimal pricing and production decisions of auto manufacturers that sell overlapping vintages of the same product simultaneously. Copeland and Hall (2011) examine how the Big Three automakers accommodate shocks to demand. Graddy and Hall (2011) compare dynamic pricing that sets one price per period based on inventory levels to pricing that also allows for third-degree price discrimination. Dana and Willams (2020) show in an oligopoly model that strong competitive forces can limit intertemporal price discrimination. Finally, Chen (2018) study profit and welfare implications of dynamic pricing techniques in a competitive setting and construct a dynamic structural model of the airline industry.

To the best of our knowledge, we are the first to empirically show that market power affects firms' ability to dynamically price.

The paper proceeds as follows. In Section 2, we describe our data and discuss the measurement of

inventory and substitute inventory in the context of automobile dealerships. In Section 3, we discuss estimation issues. In Section 4, we establish the existence of inventory-based dynamic pricing in car retailing, namely the price–inventory relationship. In Section 5, we present the main result of the paper, namely that higher market power strengthens the price–inventory relationship, and we analyze the robustness of this result. In section 6, we offer a conclusion.

2. Data and Estimation

Our data contain information on automobile transactions between January 1, 1998, and December 31, 2014, from a 30% sample of new car dealerships in the United States. A major market research firm collected the data, which include every new vehicle transaction at the dealers in the sample during the sample period. For each transaction, we observe the precise vehicle that is purchased, the price the customer paid for the vehicle, demographic information on the customer, financing information, trade-in information, dealer-added extras, and the profitability of the car and the customer to the dealership.

Before describing the different measures we construct from the data, we discuss some stylized facts about the industry to motivate the assumptions we make in our analysis. Importantly, we argue that resupply in the car retailing market is exogenous in the short to medium run.

Understanding the effect of inventory on dealer pricing depends on understanding the supply relationship between dealers and manufacturers. Technically, dealers place orders with manufacturers. Practically, however, most manufacturers have guidelines for dealers, and some manufacturers simply tell dealers which cars they will be receiving. Manufacturers force dealer ordering through bundling of cars. Dealers must take a certain number of slow-selling cars if they want an allocation of popular cars. Overall, car dealers have some input into the selection of cars and models but only a limited amount. Furthermore, their role is concentrated in the area of specifying trim levels and types of cars rather than large changes in gross quantities or models.

In interviews with car dealers and manufacturers, we found that, although dealers order frequently from manufacturers, it takes at least 45 days—and typically 90 days—for the dealer to actually receive the car. Within that time period, dealers cannot obtain additional cars from the manufacturer for delivery at that shipping date.⁵ Also, they cannot reduce their order or alter its composition.⁶ Of course, a dealer can have *expected* inventory that is a strong function of current sales by, for example, reordering every car sold. But, because of the typical 90-day lag between

order and delivery, the cars cannot be test driven or examined by customers, and neither can they be sold to customers before arrival on the lot.⁷ If a customer cannot find what the customer wants on the lot, the customer will either shop at another dealer or come back a few days later—if inventory is expected to arrive—rather than place an order. This is because shoppers tend to want to drive away with a new car on the day they shop for it. The inventory actually on the lot, therefore, retains considerable importance in pricing.

Because resupply for each dealer is exogenous in the short to medium run, so is the amount of substitute inventory in a dealer's market. Although dealers order frequently from manufacturers, the lag between ordering and delivery means that competing dealers cannot increase the amount of substitute inventory in less than 45 to 90 days.

2.1. Inventory Measurement

The first goal of this paper is to show that inventories systematically affect pricing in the car retailing industry. To establish this price–inventory relationship, we first need to define and measure inventory in our context. We provide a theoretical intuition behind this inventory-based dynamic pricing mechanism in the online appendix.

We measure inventory on the level of the interaction of make, model, model year, body type, and doors. This means that any given make and model, for example, a Honda Accord, can have different inventory levels at the same dealer, depending on whether it is the 2013 or 2014 model, whether it is manual or automatic, etc. Tracking inventory on the level of this definition is important because customers may have preferences over these attributes, and some varieties of a make and model may be in short supply while the others are not. By measuring inventory this way, we are making an assumption that customers substitute between versions of a car relatively easily (because different trim levels are substitutable in this setup). We test this assumption later in the paper.

Because our data are derived from a record of transactions, we do not have a direct measure of inventory. We do know, however, which cars were sold and how long each sat on the lot before the sale. This measure, *DaysToTurn*, allows us to derive when the car arrived on the dealer's lot. Knowing the arrival and departure dates for each car sold at each dealership allows us to construct how many cars were on the dealership's lot at any given time by “rolling back” the data. Moving from the latest sale backward, each car can be counted as part of the dealer's inventory for the number of days it was on the lot. This measure is accurate at the beginning of our sample period because all cars on a dealer's lot at that point would have been sold during our sample period of 17 years, thereby allowing us to

identify when it came on the lot. Notice, however, that our inventory measure is less accurate as we approach the last year of the sample period. This is because we only observe when cars come on the lot if they subsequently are sold during our sample period. Many cars that arrive on the lot at the end of our sample period are sold after the end of the sample. Consequently, we exclude the last 12 months of our sample from our price specifications. We choose 12 months because the days to turn for nearly all (99.9%) cars fall within this time frame. Hence, our final data set comprises car purchases for 16 years from January 1, 1998, to December 31, 2013. Figure 1 shows the inventory levels over time for a Honda dealer. We graphed the inventory levels of three typical cars over a two-month period, including when cars arrive on the lot and when they are sold.

Having measured inventory at each dealer on each day, we obtain a wide range of inventory levels (from 1 to 605 vehicles). We do not have a prior on the exact functional form that inventory should take in determining prices. One might expect that inventory would have a different relationship with prices at large versus small dealerships, and the marginal impact of a unit of inventory may be smaller for larger levels of inventory. We, therefore, considered three different methods to scale our inventory measure.

First, we considered normalizing inventory by average dealer sales volume to create a measure of inventory level relative to average sales rate. This approach proved problematic in our sample (and is, thus, not reported) because dealer inventory should not necessarily scale linearly with sales. To understand this, note that even small dealers need a certain number of cars on the lot to be able to offer variety to customers. This implies that a large dealer does not necessarily need more cars on the lot compared with a small dealer; given the same variety, the large dealer can simply choose to be resupplied more often.

Second, we considered using indicators for when a dealership's inventory is below certain percentile levels specific to the dealership. This second approach proved problematic (and is, thus, also not reported) because, given the fine granularity of our car definition, the 5th, 10th, and even 25th percentile of inventory is one for small dealerships (see the top panel of Figure 2 for a histogram of daily inventories for all dealers). This points to a larger problem, which we address next, namely that there is not much variation in inventory of a particular car for small dealerships.

We settled on a third approach, namely to restrict the sample to dealerships that sell a minimum number of cars and use the raw number of cars in inventory as our inventory measure. In this latter case, we allow for two coefficients on the marginal car, one for inventory levels below the median of 15 and one for

Figure 1. Example of Inventory Movement for Three Cars at One Honda Dealer in January and February 2013





Figure 2. Distribution of Daily Inventories (at the Inventory Car Level)



15 and above. Specifically, we restrict the sample to dealership–car combinations for which the dealership sells at least three cars per month according to our definition of a car (see the bottom panel of Figure 2 for a histogram of daily inventories for such dealership–car combinations) and then simply count the cars in inventory. In choosing this approach, we measure the average effect of an additional unit of inventory across dealers of different size.⁸ This approach leaves 9,042,402 observations.

2.2. Market Power Measurement

The second goal of this paper is to empirically show how market power affects firms’ ability to price dynamically. To do so, we need to define market power in our context. As described in the introduction, our operationalization of market power is based on substitute inventory in each dealer’s selling area.

Market power is usually defined at a more aggregate level, such as the level of a firm or brand. However, this is not necessarily a better definition. For example, suppose a firm sells two unrelated products A and B. Product A is sold in a competitive market, and the firm is a monopolist in the market for product B. We argue that, in this case, one should not define the

firm’s market power at the firm instead of the product level. In another example, suppose a firm is a monopolist in year 1. Now, assume that there are many entrants in year 2. In this case, it seems better to define market power at the firm–time period level instead of the firm level. Combining these two examples highlights why one should define market power at a more granular level in settings in which different products face different time-varying competitive forces. We believe that car retailing represents such a setting. In fact, we think that our ability to measure a source of market power at a granular level with exogenous intertemporal variation sets us apart from other papers that analyze market power and is one of the key advantages of our identification strategy.⁹

We define a dealer’s “selling area” for the purpose of measuring substitute inventory using two alternative approaches. In the first approach, we define a focal dealer’s local market as the designated market area (DMA) in which the dealer is located. DMAs are a standard measure of TV markets (e.g., Los Angeles, Santa Barbara–San Marino–San Luis Obispo, San Diego, etc.).¹⁰ In the second approach, we define a focal dealer’s local market as all dealers within a 30-mile radius of the focal dealer (see Olivares and Cachon 2009).¹¹ For each approach, we define substitute inventory for each transaction as the total number of vehicles of the same type of “car” (based on the inventory definition) that were available for sale at the time of the transaction in the focal dealer’s local market, according to our two definitions. This measure excludes the focal dealer’s own inventory.

Figure 3 presents a dealer’s own inventory levels compared with the local substitute inventory levels (based on the DMA measure) over time for the same Honda dealer from Figure 1. As the figure shows, local substitute inventory levels vary widely over time within a car–dealer combination, effectively leading to variation in market power for a dealer selling a focal car. To demonstrate the variation in market power across car–dealer combinations, for each combination, we compute the standard deviation of local substitute inventory. Figure 4 presents the distribution of these standard deviations. Roughly 10% of car–dealer combinations have no variation in local substitute inventory. These combinations are mostly cases in which there is no local substitute inventory (i.e., cases in which the dealer is the only dealer that carries a particular car in inventory). This is consistent with the fact that roughly 9% of car–dealership combinations have zero local substitute inventory on average. Figure 5 presents the distribution of the average local substitute inventory for each car–dealership combination.

We operationalize market power by classifying substitute inventory into four quartiles.¹² Quartile 1, the lowest substitute inventory, is associated with the

Figure 3. Example of Focal and Local Market Inventory Movement for Three Cars at One Honda Dealer in January and February 2013





Figure 4. Distribution of Standard Deviations of Local Substitute Inventories (at the Dealership–Car Level)




highest market power, and quartile 4 is associated with the lowest market power. Note that this definition of market power at the level of a car and local market allows for a large variation in market power over time for each car within a dealer and also allows for a large variation in market power across cars within a dealer during any given day.

2.3. Resupply Measurement

The extant literature makes predictions based on inventory levels conditional on the remaining time until a deadline. This is because the opportunity cost of selling a product changes as the deadline approaches. In our setting, there is no deadline because cars can remain on the lot indefinitely (at a cost). Instead, what changes the opportunity cost of selling a car of a particular type is that additional cars of that type are scheduled to be delivered to the dealer in the future.

Figure 5. Distribution of Average Local Substitute Inventories (at the Dealership–Car Level)



Hence, to control for the changing opportunity cost of selling a car, we measure the “days to resupply” for each car at each dealership.

The problem in defining this measure is that there are two types of car arrivals in our data. The first type is the arrival of a shipment from a manufacturer. The second type is the arrival of a car that was traded with another dealership. For both types of arrivals, the “days to turn” variable is set to zero on the car’s arrival day. We are concerned about traded vehicles because their arrival is not known in advance and should, thus, not factor into the dealer’s pricing decision in the same way as manufacturer shipments. Instead, vehicles are typically traded because a customer wants a specific car, and the dealer offers to obtain this car for the customer at another dealership in the region. According to industry participants we interviewed, such “trades” are indeed always an exchange. If the competing dealer agrees on the trade, an employee of the requesting dealership drives an agreed-upon exchange vehicle to the other dealership and brings the requested vehicle back. If the cars are of different value, dealers settle the difference at invoice prices.¹³

We use specific differences in the way that trades and regular shipments get on the dealer’s lot to identify which cars are dealer-initiated trades. In particular, we use three pieces of information: the odometer of the vehicle at the time it was sold, the number of days the vehicle was on the lot when sold, and the number of other vehicles that arrived on the dealer’s lot on the same day. The idea is as follows: if a car was not sold within the first few days of arriving on the lot, it is unlikely to be a requested trade. Among those cars that sold after only a few days on the lot, those cars with low mileage are unlikely to be requested trades. This is because a requested trade will have been driven from one dealership to the other. Also, a requested trade arrives on the dealership’s lot after having been on another dealer’s lot and perhaps having already been test driven for some time. The problem is to determine what should qualify as “low” or “high” mileage. We construct a mileage cutoff as follows. We calculate the 95th percentile of odometer mileage for each combination of car, dealer, and number of days in inventory when a car sells but only using a sample of cars for which at least three cars, according to our (very granular) inventory car definition, arrived on the lot on the same day. Because cars are traded one by one, it is highly unlikely that such a sample contains traded cars. We then define *TradeRequested* as a vehicle that is sold within four days of arriving on the lot and has an odometer reading that exceeds the 95th quantile as derived. Because every requested trade results in a received trade at the reciprocating dealer, we define a car as *TradeReceived* if it has an odometer reading that exceeds the same 95th quantile, was not

TradeRequested, and was the only car of that make that arrived on the dealership’s lot that day. Approximately 9% of vehicles are classified as *TradeRequested* and another 9% are classified as *TradeReceived* in the original sample. This matches well with industry estimates that less than 20% of sold cars are dealer trades.

We can now define *DaysToResupply* as the number of days until a vehicle of the same inventory car definition arrives, excluding vehicles that were classified as *TradeRequested* or *TradeReceived*. The distribution of *DaysToResupply* for the full and restricted data sets we use in this paper (dealership–car combinations for which the dealership sells at least three cars per month, according to our definition of a car) can be seen in Figure 6.

We use *TradeRequested* as an indicator variable. The sign of the coefficient is an empirical question. On the one hand, dealers bear additional transaction and transportation costs for requested trades and might pass these on to the buyer. On the other hand, dealers might discount trades to induce customers to wait for the trade to arrive instead of switching dealerships.

We have excluded from the data all transactions that occur 45 days or fewer before the introduction of the next model year. We omit these transactions from the data set because their resupply conditions are not

normal; instead, these prices reflect the effect of “fire sales” to clear dealer lots to prepare for the introduction of new models.

2.4. Dependent Variable


The price observed in the data set is the price that the customer pays for the vehicle, including factory-installed accessories and options and dealer-installed accessories contracted for at the time of sale that contribute to the resale value of the car.¹⁴ The *Price* variable we use as the dependent variable is this price minus the *ManufacturerRebate*, if any, given directly to the customer and minus what is known as the *TradeInOverAllowance*. *TradeInOverAllowance* is the difference between the trade-in price paid by the dealer to the customer and the estimated wholesale value of the trade-in vehicle (as booked by the dealer). We adjust for this amount to account for the possibility, for example, that dealers may offer customers a low price for the new car because they are profiting from the trade-in. Our measure of price also takes into account any variation in holdback and transportation charges.¹⁵

2.5. Controls

We include a car fixed effect for each combination of make, model, body type, transmission, displacement, doors, cylinders, and trim level.¹⁶ Although our car fixed effects control for many of the factors that contribute to the price of a car, they do not control for the factory- and dealer-installed options that vary within trim level. The price we observe covers such options, but we do not observe what options the car actually has. In order to control for price differences caused by options, we include as an explanatory variable the percentage deviation of the dealer’s cost of purchasing the particular vehicle from the manufacturer from the average cost of purchasing that car from the manufacturer in the data set. This percentage deviation, called *VehicleCost*, is positive when the specific vehicle has an unobserved option (for example a CD player) and is, therefore, relatively expensive compared with other examples of the same car (as specified). The *VehicleCost* variable also serves to control for manufacturer-to-dealer incentives.

To control for time variation in prices, we define a dummy *EndOfMonth* that equals one if the car was sold within the last five days of the month. A dummy variable *WeekEnd* specifies whether the car was purchased on a Saturday or Sunday to control for a similar, weekly effect. In addition, we introduce dummies for each month in the sample period to control for other seasonal effects and for inflation. If there are volume targets or sales on weekends, near the end of the month, or seasonally, we pick up their effect on prices with these variables.

Figure 6. Distribution of Daily Days to Resupply (at the Inventory Car Level)



We control for the number of months between the introduction of a car's model and when the vehicle was sold. This proxies for how new a car design is. Judging by the distribution of sales after car introductions, we distinguish between sales in the first four months, months 5–13, and month 14 and beyond, and we assign a dummy variable to each category.

We also control for the age, gender, income, education, occupation, race, and other demographic characteristics of buyers. We observe age and gender at the individual buyer level, and other demographic information stems from census data that we matched with the buyer's address from the transaction record. The data are on the level of a "block group," which makes up about one fourth of the area and population of a census tract. On average, block groups have about 1,100 people in them.

Finally, we control for the DMA in which the car was sold and possible unobserved dealer-specific effects (including the competitiveness of each dealer's market) through dealer fixed effects in all specifications.

2.6. Final Sample

To keep the estimation tractable in the presence of high-dimensional car and dealer fixed effects, we eliminate car types that had relatively low sales as well as small dealers from the data set. By excluding car types that sold fewer than 2,991 times over the model year nationwide and dealers who sold fewer than 615 cars over the sample period, we reduce car and dealer fixed effects by 80%, respectively. Cars with few sales over the sample period have hardly any variation in inventory levels. Hence, they are unhelpful in identifying inventory effects. A similar argument holds for dealerships with few sales. We also exclude 178 transactions with a price of more than US\$100,000. Our final data set contains 4,903,122 observations. Summary statistics for the data set are in Table 1.¹⁷

2.7. Estimation Issues

We are concerned about potential endogeneity of price and inventory levels. Our maintained assumption is that inventory changes exogenously as a result of the random arrival of customers. Instead, what could be occurring is that a dealership has a sale for some reason, and the sale (i.e., low prices) results in low inventory. To reduce the chance that we are measuring the effect of prices on inventory instead of the reverse, we measure a dealer's inventory two days before the focal transaction. Thus, transactions that occur in response to a dealership's weekend sale have as an inventory measure the dealer's inventory on the preceding Thursday. In addition, our concern is mitigated by the fact that any such endogeneity would operate in the opposite direction of the inventory effect (our results show that low inventory is associated with high prices).

Of more concern is the potential simultaneous determination of price and inventory levels resulting from a demand shock. Suppose, for example, that there is a sudden increase in consumer taste for a particular car. For example, a particularly snowy winter in a region of the country may simultaneously increase prices and run down inventories for four-wheel-drive vehicles in that region. We take two approaches to account for this potential endogeneity. Our first approach makes extensive use of car, dealer, and time fixed effects (including interactions thereof) to identify the effect of inventory on price based only on short-term variations in inventory within car and dealership combinations. This means that we rely neither on variation across dealerships nor across cars or months to identify the inventory effect. This makes it less likely that our results are due to demand shocks. Our second approach is to use exogenous plant closures as an instrument for inventory. In particular, we use plant closures that result from fires, parts shortages, floods, etc., to instrument for the dealer inventory levels of the cars produced at these plants. We discuss both approaches in more detail in the next sections.

3. Inventory-Based Dynamic Pricing

In this section, we establish the existence of a price–inventory relationship in car retailing.

3.1. Existence of the Price–Inventory Relationship in Car Retailing

Our dependent variable is *Price* as defined in the data section. In order to provide the appropriate baseline for the price of the car, we use a standard hedonic regression of log price. We work in logs because the price effect of many of the attributes of the car, such as being sold in Northern California or in a particular month, are likely to be better modeled as a percentage of the car's value than as a fixed dollar increment. We estimate the following specification:

$$\ln(\text{Price}_i) = \mathbf{X}_i\boldsymbol{\alpha} + \mathbf{D}_i\boldsymbol{\beta} + \mathbf{I}_i\boldsymbol{\gamma} + \varepsilon_i. \quad (1)$$

The X matrix is composed of transaction and car variables: car, dealer, month, and region fixed effects; car costs; and controls for whether the car was purchased at the end of a month or over a weekend. The matrix also contains an indicator for whether the buyer traded in a vehicle. The D matrix contains demographic characteristics of the buyer and the buyer's census block group. To this basic specification, we add a matrix I that contains various inventory-related explanatory variables, such as measures of inventory, days to re-supply, and a trade-requested indicator.

To estimate the effect of inventory on prices, we estimate a specification that allows us to test the standard prediction of dynamic pricing models under inventory, namely that prices should decrease in inventory,

Table 1. Summary Statistics

Variable	Mean	Standard deviation	Minimum	Maximum	N
<i>Price</i>	25,658.04	7,817.37	5,990	100,000	4,903,122
<i>Inventory</i>	29.46	35.65	1	605	4,903,122
<i>DaysToResupply</i>	6.92	14.11	1	996	4,903,122
<i>LocalSubstituteInventory</i>	120.49	181.69	0	2461	4,903,122
<i>TradedCar</i>	0.09	0.29	0	1	4,903,122
<i>Tradein</i>	0.46	0.5	0	1	4,903,122
<i>%Black</i>	0.07	0.15	0	1	4,903,122
<i>%Hispanic</i>	0.12	0.19	0	1	4,903,122
<i>%Asian</i>	0.05	0.09	0	1	4,903,122
<i>Female</i>	0.41	0.49	0	1	4,903,122
<i>Income</i>	60,213.02	25,548.91	0	200,001	4,903,122
<i>Income</i> ²	4,278,354,877	4,074,203,432	0	40,000,401,408	4,903,122
<i>%LessHighSchool</i>	0.14	0.12	0	1	4,903,122
<i>%CollegeGrad</i>	0.39	0.19	0	1	4,903,122
<i>%Management</i>	0.17	0.08	0	1	4,903,122
<i>%HProfessional</i>	0.23	0.1	0	1	4,903,122
<i>%Health</i>	0.02	0.02	0	1	4,903,122
<i>%Protective</i>	0.02	0.02	0	1	4,903,122
<i>%Food</i>	0.04	0.03	0	1	4,903,122
<i>%Maintenance</i>	0.03	0.03	0	1	4,903,122
<i>%Housework</i>	0.03	0.02	0	1	4,903,122
<i>%Sales</i>	0.12	0.05	0	1	4,903,122
<i>%Admin</i>	0.16	0.05	0	1	4,903,122
<i>%Construction</i>	0.05	0.04	0	1	4,903,122
<i>%Repair</i>	0.04	0.03	0	1	4,903,122
<i>%Production</i>	0.06	0.05	0	1	4,903,122
<i>%Transportation</i>	0.05	0.04	0	1	4,903,122
<i>MedianHHSize</i>	2.72	0.53	0	8.93	4,903,122
<i>MedianHouseValue</i>	181,360	123,763	0	1,000,001	4,903,122
<i>VehPerHoushold</i>	1.83	0.38	0	7	4,903,122
<i>%HouseOwnership</i>	0.74	0.23	0	1	4,903,122
<i>%Vacant</i>	0.06	0.07	0	1	4,903,122
<i>TravelTime</i>	27.7	6.79	0	200	4,903,122
<i>%Unemployed</i>	0.04	0.04	0	1	4,903,122
<i>%BadEnglish</i>	0.04	0.08	0	1	4,903,122
<i>%Poverty</i>	0.08	0.08	0	1	4,903,122
<i>CustomerAge</i>	46.56	14.61	16	110	4,903,122
<i>Age > 64</i>	0.12	0.33	0	1	4,903,122
<i>VehicleCost</i>	0	0.06	-0.78	1.16	4,903,122
<i>Model Age 5–13 Months</i>	0.69	0.46	0	1	4,903,122
<i>Model Age >14 Months</i>	0.13	0.33	0	1	4,903,122
<i>Weekend</i>	0.3	0.46	0	1	4,903,122
<i>EndOfMonth</i>	0.25	0.43	0	1	4,903,122
<i>EndOfYear</i>	0.03	0.17	0	1	4,903,122

controlling for days to resupply. Because one additional car in inventory may have a different effect on price if inventory levels are low versus high, we include the inventory variable as a two-part spline in our specification (in the online appendix, we show that this fits the data well). In particular, we estimate a different inventory coefficient for below- and above-median inventory levels (the median is 15) while controlling for days to resupply. This initial specification includes both car and dealer fixed effects. We include dealer fixed effects to be able to identify the price–inventory relationship within and not across dealers. If we did not include dealer fixed effects, we would be concerned that the hypothesized negative price–inventory relationship could be due to large dealers that simultaneously have

higher absolute inventory levels and lower prices because they are more cost-efficient than small dealers.

Column (1) of Table 2 reports the results of estimating this specification. Both inventory coefficients have the hypothesized negative sign. For below-median inventory levels (14 and fewer cars), one additional car in inventory is associated with a price that is lower by 0.039% (see variable *Inventory (1–14)*). For above-median inventory levels (15 and more cars), one additional car in inventory is associated with a price that is lower by 0.0057% (see variable *Inventory (15+)*). An increase in inventory from 1 to 37 cars (a one-standard-deviation increase) is associated with a 0.64% reduction in average price. This corresponds to \$164 or 36.5% of the average dealer gross margin on a vehicle

Table 2. Basic Result: Price Effects of Inventory^a

	Fixed effects			
	Car, dealer, month	Car × dealer month	Car × dealer, month × segment × DMA	Car, dealer, week × subsegment × DMA
<i>Inventory (1–14)</i>	–0.039** (0.00087)	–0.046** (0.0011)	–0.038** (0.0012)	–0.027** (0.001)
<i>Inventory (15+)</i>	–0.0057** (0.00013)	–0.005** (0.00022)	–0.0033** (0.00024)	–0.0037** (0.00016)
<i>DaysToResupply</i>	–0.0022** (0.00023)	–0.00097** (0.0003)	–0.0001 (0.0003)	–0.0009 (0.00026)
<i>TradedCar</i>	–0.18** (0.0087)	–0.22** (0.0092)	–0.23** (0.0091)	–0.19** (0.009)
<i>Tradein</i>	2.6** (0.006)	2.6** (0.0061)	2.5** (0.0061)	2.6** (0.0062)
<i>VehicleCost</i>	84** (0.092)	86** (0.099)	86** (0.098)	85** (0.093)
<i>Model Age 5–13 Months</i>	0.047** (0.011)	0.045** (0.011)	–0.061** (0.013)	0.013 (0.015)
<i>Model Age >14 Months</i>	–0.043* (0.022)	–0.028 (0.022)	–0.089** (0.025)	–0.014 (0.029)
<i>Weekend</i>	0.074** (0.0062)	0.085** (0.0063)	0.082** (0.0062)	0.058** (0.0065)
<i>EndOfMonth</i>	–0.52** (0.0068)	–0.5** (0.0068)	–0.49** (0.0068)	–0.17** (0.015)
<i>EndOfYear</i>	–0.19** (0.02)	–0.16** (0.02)	–0.16** (0.02)	–0.098* (0.04)
Observations	4,903,122	4,903,122	4,903,122	4,903,122
Adjusted R ²	0.956	0.959	0.960	0.958

^aUnreported are the constant term, car, dealer, and month fixed effects (column (1)); car × dealer and monthly fixed effects (column (2)); car, dealer, and month × segment × DMA fixed effects (column (3)); and car, dealer, week × subsegment × DMA fixed effects (column (4)); and the demographic variables reported in Table 1. All coefficients are multiplied by 100. Robust standard errors in parentheses.

*Significant at 10%; *significant at 5%; **significant at 1%.

in our sample. An increase in inventory by one standard deviation when the inventory for that car is already high has a smaller effect. For example, an increase in inventory from 15 to 51 cars is associated with a 0.205% lower average price. This corresponds to 12% of the average dealer gross margin.

The findings on the effect of inventory levels are consistent with the comparative static hypothesized by dynamic inventory models with inventory. Controlling for the time until a new shipment arrives, prices decrease as there are more cars in inventory.

The “days to resupply” control has a negative coefficient. A decrease in days to resupply by one day is associated with a 0.0022% increase in average price. This result is in line with Lin and Sibdari (2009), who show that, under competition, the optimal price for a product need not be nondecreasing in time to go.¹⁸

Highlighting some other results, we find that consumers pay a lower price (0.18%) for a vehicle that

was requested from another dealership (*TradedCar*). This is consistent with dealers discounting trades to induce customers to wait for the trade to arrive instead of switching dealerships. Cars that are sold at the end of the month (*EndOfMonth*), when salespeople are trying to meet sales quotas, sell for on average 0.52% lower prices. Demographic variables are unreported in Table 2 but have the expected sign. For example, women and minorities pay slightly more for a car, as do consumers who live in neighborhoods with a higher percentage of residents who have less than a high school education.¹⁹

3.2. Endogeneity Concerns

We would like to make sure that the estimated price–inventory relationship is not the result of a potential endogeneity of prices and inventory levels because of demand shocks. We use a sequence of fixed effects to address the potential endogeneity of price and

inventory. In the online appendix, we also present an instrumental-variables approach to estimate the effect of inventory on price levels.

In the next two specifications, we repeat the basic specification of column (1) of Table 2 with different sets of fixed effects to address the potential endogeneity of price and inventory resulting from common demand shocks. We focus on the demand shocks we feel are most plausible for the market we are studying.

So far, we have included a fixed effect for each month in our sample, for each car (with the preceding detailed definition), and for each dealer. Our first alternative specification accounts for the possibility that there are car–dealership interactions that may be responsible for our result. For example, suppose that 7 Series BMWs are particularly popular in Beverly Hills. This leads to high prices and low inventory levels at the Beverly Hills BMW dealer and, thus, forms an alternative explanation for why we find that low inventory levels may be associated with higher prices.²⁰ To rule out this alternative explanation, we repeat the specification in column (1) of Table 2 with interacted car and dealer instead of separate car and dealer fixed effects. This absorbs the mean price level for each car at each dealership separately; the price–inventory relationship is, thus, only identified from inventory fluctuations over time within car–dealer combinations. The results in column (2) of Table 2 are very similar to those of column (1): for below-median inventory levels (14 and fewer cars), one additional car in inventory is associated with a price that is lower by 0.046% (see variable *Inventory* (1–14)). For above-median inventory levels (15 and more cars), one additional car in inventory is associated with a price that is lower by 0.005% (see variable *Inventory* (15+)). Both coefficients remain precisely estimated despite a substantial decrease in degrees of freedom: although the specification in column (1) of Table 2 contains 6,705 car fixed effects and 2,740 dealer fixed effects, the specification in column (2) contains 359,236 car × dealer fixed effects.

Our second alternative specification accounts for the possibility that demand shocks are short-lived and local. So far, our monthly fixed effects absorb the price effect of short-term demand shocks but only if these affect all vehicle segments in all markets equally. This may not be a good assumption: for example, suppose that a particularly snowy January in the California Sierras increases demand for SUVs for the rest of the winter in the Sacramento area (but not in Southern California), thus simultaneously causing high prices and low inventories for the SUV segment in Sacramento dealerships for that quarter. To rule out this alternative explanation, in column (3) of Table 2, we repeat the specification of column (2) of Table 2, expanding the month fixed effects to month–local area–vehicle segment fixed effects. The local areas are defined as DMAs. This set of fixed

effects absorbs demand shocks specific to a segment (e.g., compact, SUV, pickup trucks, etc.) in a local market for a particular month. This specification contains 359,236 car × dealer fixed effects and 68,099 month × segment × DMA fixed effects (see column (3) of Table 2). We find that, for below-median inventory levels (14 and fewer cars), one additional car in inventory is associated with a price that is lower by 0.038% (see variable *Inventory* (1–14)). For above-median inventory levels (15 and more cars), one additional car in inventory is associated with a price that is lower by 0.0033% (see variable *Inventory* (15+)). Both variables remain precisely estimated. In summary, the negative price–inventory relationship seems robust across specifications that account for a variety of unobserved demand shocks as possible sources of causation. The days to resupply variables are negative and significant in some but not other fixed effects specifications.

We have also estimated the price–inventory relationship with fixed effects that absorb average weekly prices on a subsegment–DMA level. Specifically, we repeated the specification in column (1) of Table 2 with car fixed effects (6,705), dealership fixed effects (2,740), and week × subsegment × DMA fixed effects (332,465). The results are reported in column (4) of Table 2. The inventory level results continue to hold: for below-median inventory levels (14 and fewer cars), one additional car in inventory is associated with a price that is lower by 0.027% (see variable *Inventory* (1–14)). For above-median inventory levels (15 and more cars), one additional car in inventory is associated with a price that is lower by 0.0037% (see variable *Inventory* (15+)).

In this section, we have empirically shown that inventories systematically affect pricing in the car retailing industry. For tractability and because the results are similar across fixed effects specifications, we use the specification in column (2) of Table 2 as the basis for further analysis. Next, we use this result and show that the slope of the price–inventory relationship is significantly steeper when dealers find themselves in a situation of high rather than low market power.

4. Market Power and the Price–Inventory Relationship

As described in the introduction, there are two (related) sources of market power in auto retailing. First, a dealer’s market power depends on the number of competing dealers within the selling area. Second, holding constant the number of competing dealers, a dealer’s market power also varies with the quantity of substitute inventory available for sale by competing dealers. The number of competing dealers is stable in the medium run. In contrast, the amount of substitute inventory is quite volatile because it is subject to demand shocks. In this section, we empirically show that a dealer’s ability

to adjust prices in response to inventory depends on the second source of market power, that is, the quantity of substitute inventory in the selling area. In particular, we show that the slope of the price–inventory relationship (higher inventory lowers prices) is significantly steeper when dealers find themselves in a situation of high rather than low market power.

4.1. The Slope of the Price–Inventory Relationship

To estimate the effect of market power on the price–inventory relationship, we determine for each vehicle for sale at a dealer the substitute inventory for that vehicle in the focal dealer’s selling area. We use two different definitions for the local selling area of each dealer. First, we define the local market of a focal dealer as all dealers in the focal dealer’s DMA. Second, we define the local market of a focal dealer as all other dealers located within a 30-mile radius.²¹ In both specifications, we omit the focal dealer’s inventory from the sum of total substitute inventory in the local selling area. For each definition, we categorize substitute inventory into quartiles. The core quantity of interest is the coefficient on the interaction of the market power quartiles with our two-part inventory spline. Specifically, we estimate the following regression:

$$\ln(\text{Price}_i) = X_i\alpha + D_i\beta + \text{Inv}_i \times M_i\theta + M_i\delta + I_i\gamma + \varepsilon_i, \quad (2)$$

where the X matrix is composed of transaction and car variables: car, dealer, month, and region fixed effects; car costs; and controls for whether the car was purchased at the end of a month or over a weekend. The matrix also contains an indicator for whether the buyer traded in a vehicle. The D matrix contains demographic characteristics of the buyer and the buyer’s census block group. Inv_i contains the inventory spline and is interacted with the M matrix, which contains the local inventory quartile dummies. The θ coefficients are the coefficients of interest, which allow us to examine the slope of the price inventory relationship for different levels of market power. In addition, because market power may not only affect the price–inventory relationship but also price levels, we control for substitute inventory quartiles, M , directly. Finally, the I matrix contains the inventory-related controls, such as days to resupply, and a trade-requested indicator.

We test two predictions using this equation. First is that higher levels of substitute inventory are associated with lower price levels. That is, we predict that the δ vector is decreasing in substitute inventory. Our second and key prediction is that the slope of the price–inventory relationship is smaller in magnitude the more inventory competing dealers have of the same type. That is, the θ vectors are decreasing in substitute inventory.

We report the results of two specifications in Table 3, one for each definition of the local market of the dealer. Column (1) reports the results using DMA to define the local market. Low levels of substitute inventory (quartile 1) proxy for high market power, and high levels of substitute inventory proxy for low market power. Consistent with our first prediction, the price levels are decreasing in substitute inventory (note the monotonic decreasing relationship for variables (see variable *Local qX*)). Specifically, not only are the different degrees of market power different from the omitted category, *Local q1*, which is the highest market power, but also, they are statistically different from each other. Moving from a situation of high market power (local q1) to low market power (local q4) lowers transaction prices by 1% or 57% of the average dealer margin.

Our second and key prediction is tested using the interaction between the inventory spline and the market power quartiles. We find that more substitute inventory in the DMA leads to a weaker (less negative) price–inventory relationship. The interaction coefficients are statistically different from each other (except for the interactions of local q3 and q4 with below-median inventory levels, which are only marginally significant).

When there is a shortage of substitute inventory (quartile 1), a dealership moving from a situation of (own) inventory shortage to a median (own) inventory level lowers transaction prices by about 0.57% ceteris paribus, corresponding to 32.5% of dealers’ average per-vehicle profit margin or \$145.6 on the average car. Conversely, when there is ample substitute inventory (quartile 4), moving from inventory shortage to a median inventory level lowers transaction prices by about 0.35% ceteris paribus, corresponding to \$90.9 or 20.2% of dealers’ average per-vehicle profit margin. For quartiles 2 and 3, we find intermediate effects at 0.51% and 0.43%, respectively. Figure 7 summarizes the effect sizes graphically and illustrates that the slopes are smaller in magnitude the more inventory competing dealers have of the same type.²² Overall, as hypothesized, dynamic pricing is more pronounced when dealers have more market power.

The results so far confirm our predictions: First, higher levels of substitute inventory are associated with lower prices. Second, the slope of the price–inventory relationship is smaller in magnitude the more inventory competing retailers have of the same type. These results are robust to the definition of local selling area as a DMA or a 30-mile radius around each dealer. We examine additional robustness of our inventory measures in Section 4.3.

4.2. Financing and Insurance Margins

In addition to the margin on the sale of the vehicle and the trade-in, dealers and salespeople earn a margin from car F&I margins. In this section, we test

Table 3. Main Results: Local and Focal Inventory Effects on Prices^a

Dependent variable: <i>ln(price)</i>	(1) Local inventory defined at DMA level	(2) Local inventory defined at 30-mile radius
<i>Local q2</i>	-0.27** (0.024)	-0.19** (0.024)
<i>Local q3</i>	-0.64** (0.031)	-0.6** (0.031)
<i>Local q4</i>	-1** (0.045)	-0.98** (0.045)
<i>Local q1 × Inventory (1–14)</i>	-0.043** (0.0018)	-0.041** (0.0018)
<i>Local q2 × Inventory (1–14)</i>	-0.039** (0.0018)	-0.043** (0.0018)
<i>Local q3 × Inventory (1–14)</i>	-0.033** (0.0021)	-0.029** (0.0021)
<i>Local q4 × Inventory (1–14)</i>	-0.027** (0.0029)	-0.025** (0.003)
<i>Local q1 × Inventory (15+)</i>	-0.0084** (0.00049)	-0.0082** (0.00053)
<i>Local q2 × Inventory (15+)</i>	-0.0054** (0.00051)	-0.0065** (0.00052)
<i>Local q3 × Inventory (15+)</i>	-0.0043** (0.00041)	-0.0054** (0.00039)
<i>Local q4 × Inventory (15+)</i>	-0.0031** (0.00026)	-0.0031** (0.00026)
<i>DaysToResupply</i>	-0.0009** (0.0003)	-0.00083** (0.0003)
<i>TradedCar</i>	-0.22** (0.0092)	-0.22** (0.0092)
<i>Tradein</i>	2.6** (0.0061)	2.6** (0.0061)
<i>VehicleCost</i>	86** (0.099)	86** (0.099)
<i>Model Age 5–13 Months</i>	0.069** (0.011)	0.066** (0.011)
<i>Model Age > 14 Months</i>	-0.019 (0.022)	-0.02 (0.022)
<i>Weekend</i>	0.085** (0.0063)	0.088** (0.0063)
<i>EndOfMonth</i>	-0.49** (0.0068)	-0.49** (0.0068)
<i>EndOfYear</i>	-0.17** (0.02)	-0.17** (0.02)
Observations	4,903,122	4,903,122
Adjusted R ²	0.959	0.959

^aUnreported are the constant term, car×dealer, monthly fixed effect, and the demographic variables reported in Table 1. All coefficients are multiplied by 100. Robust standard errors in parentheses.

[†]Significant at 10%; *significant at 5%; **significant at 1%.

whether F&I margins, another component of price, are also affected by inventories. During a new car sale, after the customer agrees on a price with the salesperson, the customer is then sent to the F&I specialist, who—in the process of doing the paperwork with the customer—offers financing, insurance, and service products. Specifically, the F&I measure we observe captures the total profit made on (a) the sale of accident and health insurance, (b) the sale of credit life insurance, (c) the sale of service contracts, and (d) by marking up the finance or lease APR rate. F&I charges can also be negotiated, and therefore, we examine whether the F&I margins are also affected by inventories. Our sample is limited to those transactions in which F&I sales took place. We use F&I margins as the dependent variable in both the basic specification (Equation (1)) and the market power specification (Equation (2)).


Column (1) of Table 4 reports the results of the basic specification. The coefficient for below-median levels of inventory (see variable *Inventory (1–14)*) is -0.005% suggesting that a dealership moving from a situation of shortage of a particular car (one car in inventory) to a median inventory level of cars (15) lowers F&I margins by about 0.065%. The average F&I margin in the data is \$858, suggesting that, although the effects are statistically significant, 0.065% is of a relatively small economic magnitude of roughly 56 cents. The coefficient on above-median levels of inventory is also significantly different than zero at -0.0003%.

In column (2) of Table 4, we explore whether this F&I-inventory relationship also depends on market power. Again, for below-median levels of own inventory (see interactions with variable *Inventory (1–14)*), the slope is steeper when dealers have more market power. However, this result only partially carries over for above-median levels of own inventory. In particular, the pattern of decreasing margin as the level of competition increases holds only for quartiles 2–4, and the quartile 1 coefficient is not statistically different than zero. Overall, dealers’ ability to dynamically price F&I options is (slightly) weakened as the quantity of substitute inventory increases.

4.3. Robustness

We now explore the robustness of the effect of market power on inventory-based dynamic pricing. First, we test whether the results are robust to the level on which inventory is measured. In particular, we want to make sure that our estimates are not biased by the definition of a car we use for constructing our inventory measure (see Section 4.3.1). Second, we test whether the results depend on the way we measure market power (see Section 4.3.2). Third, we examine whether the results depend on the ability of consumers to access information about substitute inventories

Figure 7. (Color online) Market Power Price–Inventory Relationship Slopes



in local dealerships (see Section 4.3.3). We use the DMA-based definitions of market power in our robustness tests.

4.3.1. Is Our Inventory Measure Too Broadly Defined?

We have so far measured inventory based on a particular definition of a car. This may lead us to overestimate the effect of inventory on prices if consumers do not consider cars for which we count inventory jointly to be close substitutes. Because we are using substitute inventory as a proxy for market power, we want to make sure our results hold for a more granular level of inventory.

We analyze whether our inventory definition affects our results by defining cars at a more granular level. We redefine our inventory measures at the level of the interaction of make, model, model year, body type, doors, transmission type, and trim level. Note that this change affects both the measure of each dealer’s focal inventory as well as the measure of substitute local inventory. The results in column (1) of Table 5 show that the monotone relationship between price and market power persists as can be seen by the coefficients for *local qX* variables. However, the results on market power and the slope of the price–inventory relationship remain only for the above-median level of the focal dealer’s inventory. For below-median focal inventory, the hypothesized interaction is not present. Figure 8 illustrates these results.

In summary, most of our results are robust to a change in the level at which we measure inventory. When we use the more granular level of inventory, the differences in the slopes occur because of above-median focal inventory but not below-median inventory. One interpretation is that consumers are willing

to substitute very similar cars (which, in our narrower inventory definition, are defined as a different car) when inventories are low.

4.3.2. Do Our Results Depend on the Granularity of Our Market Power Measure?

So far, we have measured market power based on the quartile of substitute cars available in the local market. Here, we test whether our results are robust to the granularity of this market power definition. To do so, we run similar specifications to Equation (2) but with median and decile splits of the substitute inventory. For clarity of presentation, we present the results in a figure.

Figure 9 illustrates the results. For the median split results, we observe a similar pattern to the one we have seen so far. For the deciles, we generally observe the same monotone pattern except for two main differences. First, there is a reversal between the two lowest deciles, d1 and d2, such that the d2 slope is the steepest. Interestingly, d1 includes only cases in which substitute inventory is exactly zero, that is, the focal dealer is a monopoly in the local market, and d2 includes cases in which substitute inventory contains one to nine cars. Second, the slopes of the two highest deciles, d9 and d10, yield the pattern we expect (d10 interaction coefficient is smallest in magnitude) only for high levels of focal inventory (around 105 cars). Note that the result that higher levels of substitute inventory are associated with higher prices (which are not presented in the chart) is robust to the median and decile definition (again, except for d1 and d2, which are not statistically different from each other). In conclusion, our baseline results are robust to different degrees of granularity of our market power definition.

Table 4. Financing and Insurance Results^a

Dependent variable: $\ln(F&I)$	(1)	(2)
<i>Inventory (1–14)</i>	–0.005** (0.00048)	
<i>Inventory (15+)</i>	–0.0003** (0.00008)	
<i>DaysToResupply</i>	–0.00003 (0.0001)	–0.00003 (0.0001)
<i>local q2</i>		–0.03** (0.01)
<i>local q3</i>		–0.07** (0.013)
<i>local q4</i>		–0.095** (0.02)
<i>local q1 × Inventory (1–14)</i>		–0.007** (0.00072)
<i>local q2 × Inventory (1–14)</i>		–0.004** (0.00074)
<i>local q3 × Inventory (1–14)</i>		–0.0025** (0.00089)
<i>local q4 × Inventory (1–14)</i>		–0.0022 (0.0014)
<i>local q1 × Inventory (15+)</i>		–0.0001 (0.00017)
<i>local q2 × Inventory (15+)</i>		–0.00097** (0.00018)
<i>local q3 × Inventory (15+)</i>		–0.00055** (0.00015)
<i>local q4 × Inventory (15+)</i>		–0.00018 ⁺ (0.0001)
<i>TradedCar</i>	–0.11** (0.0043)	–0.11** (0.0043)
<i>Tradein</i>	0.098** (0.0024)	0.098** (0.0024)
<i>VehicleCost</i>	0.26** (0.025)	0.26** (0.025)
<i>Model Age 5–13 Months</i>	–0.029** (0.0047)	–0.027** (0.0047)
<i>Model Age > 14 Months</i>	–0.033** (0.0089)	–0.032** (0.0089)
<i>Weekend</i>	0.064** (0.0025)	0.064** (0.0025)
<i>EndOfMonth</i>	–0.058** (0.0028)	–0.057** (0.0028)
<i>EndOfYear</i>	0.0038 (0.0082)	0.0037 (0.0082)
Observations	2,758,335	2,758,335
Adjusted R ²	0.211	0.212

^aUnreported are the constant term, car×dealer, monthly fixed effects, and the demographic variables reported in Table 1. All coefficients are multiplied by 100. Robust standard errors in parentheses.

⁺Significant at 10%; *significant at 5%; **significant at 1%.

4.3.3. Do Consumers Need Access to Information About Substitute Inventories? Consumers have always had the ability to physically visit other dealers to learn about their inventory. Such search, however, is quite costly, in particular at dealers who are not in the consumer’s close vicinity. Starting in 1999—as automobile manufacturers and dealers started adding inventory features into their websites—these search costs started decreasing. Inventory listing on websites allowed consumers to easily observe the dealers’ inventories before negotiating for prices. AutoNation dealerships started posting inventory information in July 1999, Chrysler in 2000, Chevrolet in 2001, Ford in 2002, GMC in 2003, Toyota in 2006, and all other manufacturers in 2007. In this section, we investigate whether our results hold even when consumer search for substitute inventories is costly.

To investigate, we split our sample to “online information” and “off-line information” periods based on whether inventory information could have potentially been obtained online, using the timing of when inventory information was made available to consumers. The results are reported in columns (2) and (3) of Table 5. We examine the coefficients of the interactions between focal inventory and market power. The online results replicate our existing findings regarding the slope of the price–inventory relationship. However, for the off-line results, for each of the inventory splines, we do not find a monotone relationship between market power and the effect on price. In fact, for each of the two inventory splines, the coefficients are not statistically different from one another (except for the coefficient of the third quartile for above-median inventory levels, which is different than the second and fourth quartile coefficients). In other words, when it is costly for consumers to observe other dealers’ inventories, the dealer’s ability to adjust pricing in response to inventory does not depend on the quantity of substitute inventory in the dealer’s selling area. Our empirical results seem to depend on consumers’ ability to easily observe competitive dealer inventories.

It is beyond the scope of this paper to build a theory that links inventory information to the inventory–price relationship (this would require an equilibrium model that associates how dealers would react to what consumers know). However, we can use theory to form a hypothesis on how inventory information might affect price levels. A class of bargaining theoretic models investigates the relationship between information asymmetries among bargaining parties and the division of surplus obtained in a negotiation (see section 5.1 in Busse et al. (2006) for the relevant literature). In these models, reducing the information asymmetry of a party allows that party to obtain a larger share of the surplus in the negotiation. In our setting, consumers are initially uninformed about inventories, and they


Table 5. Robustness: Inventory Definition and Inventory Information^a

Dependent variable: $\ln(\text{price})$	(1) Narrower inventory definition	(2) Online information	(3) Off-line information
<i>local q2</i>	-0.18** (0.021)	-0.35** (0.028)	-0.13** (0.043)
<i>local q3</i>	-0.43** (0.028)	-0.82** (0.037)	-0.23** (0.06)
<i>local q4</i>	-0.76** (0.044)	-1.2** (0.051)	-0.66** (0.093)
<i>local q1 × Inventory (1–14)</i>	-0.054** (0.0036)	-0.049** (0.0021)	-0.033** (0.0036)
<i>local q2 × Inventory (1–14)</i>	-0.065** (0.0033)	-0.039** (0.0021)	-0.04** (0.0033)
<i>local q3 × Inventory (1–14)</i>	-0.065** (0.0037)	-0.03** (0.0025)	-0.041** (0.0041)
<i>local q4 × Inventory (1–14)</i>	-0.07** (0.006)	-0.025** (0.0034)	-0.033** (0.0063)
<i>local q1 × Inventory (15+)</i>	-0.011** (0.00086)	-0.009** (0.00054)	-0.0056** (0.0013)
<i>local q2 × Inventory (15+)</i>	-0.0074** (0.00097)	-0.0067** (0.0006)	-0.0045** (0.001)
<i>local q3 × Inventory (15+)</i>	-0.0062** (0.0007)	-0.0044** (0.00047)	-0.007** (0.00085)
<i>local q4 × Inventory (15+)</i>	-0.0034** (0.00041)	-0.0036** (0.0003)	-0.0033** (0.0006)
<i>DaysToResupply</i>	-0.00025 (0.00016)	-0.00057 ⁺ (0.00034)	-0.00055 (0.00065)
<i>TradedCar</i>	-0.23** (0.0092)	-0.21** (0.011)	-0.25** (0.017)
<i>Tradein</i>	2.6** (0.0061)	2.2** (0.0071)	3.5** (0.012)
<i>VehicleCost</i>	86** (0.099)	87** (0.12)	84** (0.19)
<i>Model Age 5–13 Months</i>	0.058** (0.011)	0.042** (0.013)	0.065** (0.023)
<i>Model Age > 14 Months</i>	-0.021 (0.022)	-0.017 (0.026)	-0.056 (0.043)
<i>Weekend</i>	0.081** (0.0063)	0.088** (0.0072)	0.081** (0.012)
<i>EndOfMonth</i>	-0.5** (0.0068)	-0.54** (0.008)	-0.36** (0.013)
<i>EndOfYear</i>	-0.17** (0.02)	-0.19** (0.024)	-0.15** (0.037)
Observations	4,903,122	3,584,401	1,381,721
Adjusted R^2	0.959	0.958	0.960

^aUnreported are the constant term, car×dealer, monthly fixed effects, and the demographic variables reported in Table 1. For narrower inventory definition, the median is eight cars, and the splines are adjusted accordingly. All coefficients are multiplied by 100. Robust standard errors in parentheses.

⁺Significant at 10%; *significant at 5%; **significant at 1%.

Figure 8. (Color online) Robustness: More Granular Definition of a Car



are only revealed for each dealership once consumers visit that particular dealership. We can interpret adding inventory features into websites as reducing the information asymmetry between dealers and these consumers. Therefore, we hypothesize that consumers' ability to observe inventories results in lower prices.


Our results support this hypothesis: comparing the coefficients for the levels of local inventory across the two columns, we find that, for each quartile of market power, price reductions in the online condition are larger compared with those in the off-line condition.

5. Conclusion

In this paper, we first demonstrate that the new vehicle market in the United States is subject to inventory-based dynamic pricing. We present evidence that local dealer inventory has a statistically and economically significant effect on the prices at which new cars are sold. A dealership moving from a situation of shortage to a median inventory level lowers transaction prices by about 0.51% *ceteris paribus*, corresponding to 29% of average dealer margins or \$132 on the average car. We do not find consistent evidence on the relationship between resupply times and transaction prices.

Our second and principal goal is to investigate how market power affects firms' ability to dynamically price. To do so, we leverage exogenous inventory fluctuation as a measure of market power and then explore how the price–inventory relationship varies with said market power. As hypothesized, we find that lower market power (as measured by higher levels of substitute inventory) is associated with lower average prices and that prices increase with market power. In addition, we find that the degree of market power also changes the price–inventory relationship at dealers. In particular, the slope of the price–inventory relationship is smaller in magnitude the more competitive the market. When there is a shortage of substitute inventory (quartile 1), a dealership moving from a situation of (own) inventory shortage to a median (own) inventory level lowers transaction prices by about 0.57% *ceteris paribus*, corresponding to 32.5% of

Figure 9. (Color online) Robustness: Different Definitions of Market Power



dealers' average per vehicle profit margin or \$145.6 on the average car. Conversely, when there is ample substitute inventory (quartile 4), moving from inventory shortage to a median inventory level lowers transaction prices by about 0.35% *ceteris paribus*, corresponding to \$90.9 or 20.2% of dealers' average per vehicle profit margin. For quartiles 2 and 3, we find intermediate effects at 0.51% and 0.43%, respectively. Overall, dynamic pricing is more pronounced when dealers have more market power. We also find a similar relationship between financing and insurance margins and inventory.

To our knowledge, we are the first to empirically show that market power affects firms' ability to dynamically price. In addition, our paper has implications for our understanding of dealer behavior and consumer and manufacturer strategies.

Our basic results on the price–inventory relationship shed light on why most dealers use a negotiated price instead of a fixed price strategy. Consumer advocates argue that this practice allows dealers to discriminate between consumers with a different willingness to pay or ability to bargain. Indeed, many consumers find “haggling” stressful: for example, according to a 2016 study, “More than three in five Americans (61%) feel like they’re taken advantage of at least some of the time when shopping at a car dealership.”²³ Our paper suggests that there is another reason why dealers offer cars at varying prices to shoppers: dealers can incorporate the latest information on inventory levels into the offered price. As a result, the opportunity cost to the dealer of selling a car—and, therefore, the transaction price—likely varies across two customers who purchase the same car on different days even if their willingness to pay and their bargaining ability are similar.

We also believe that consumers can learn from our results on dynamic pricing and market power. First, car-buying advice often suggests that consumers should “shop around.”²⁴ Our paper shows that one benefit of doing so is to uncover dealers with high inventory positions, which generally makes these dealers willing to accept lower price offers. Second, our paper suggests that a dealer with low inventory will not necessarily offer high prices. The dealer's ability to extract scarcity rents depends on the available substitute inventory in the local market. Therefore, in evaluating a dealer's inventory position to determine whether it is likely to indicate low or high prices, consumers benefit from knowing the inventory in the local market.

Finally, our results have implications for manufacturer strategies. First, some industry observers have commented that the dealer networks of U.S. manufacturers are too big and, therefore, depress dealer margins.²⁵ Our results suggest that increasing dealer margins would take more than thinning out the dealer networks. Manufacturers also need to

manage substitute inventory: we have shown that large substitute inventory, *even at the DMA level*, not only decreases average prices, but it also harms a dealer's ability to take advantage of scarce inventory to increase margins.

Second, an argument for guaranteeing higher margins to dealers is that it allows them to invest in customer service (loaners, showrooms, valet service, etc.) to improve the customer experience. Our paper shows that managing substitute inventory may be one lever to achieve this. However, we also show that lower substitute inventory goes hand in hand with a dealer's ability to dynamically price in response to demand shocks. This means that customers who come in at different times may pay very different prices for the same vehicle, a shopping environment that consumers are likely to perceive as a haggle environment. This is likely to be at odds with what consumers perceive as a high-service setting. Therefore, manufacturers may need additional (contractual) levers to implement a high-service shopping experience driven by high retail margins together with a low-haggle approach.

Acknowledgments

The authors are grateful for helpful comments from Eva Ascarza, Severin Borenstein, Dennis Carlton, Jose Silva, K. Sudhir, Candi Yano, Jan Van Miegen, and Martin Lariviere as well as participants at the Junior Faculty in Marketing Science Forum, the National Bureau of Economic Research Industrial Organization (NBER IO) Program Meeting, and the Marketing in Israel Conference and seminar participants at Harvard Business School, Interdisciplinary Center Herzliya, Northwestern University, Stanford University, University of California Berkeley, the University of Chicago, the University of East Anglia, and the University of Rochester. The authors are especially grateful to Meghan Busse for extensive comments and to Thomas Hubbard for the suggestion to look into inventories.

Endnotes

¹ Even if interdealer vehicle trades mean that supply is not absolutely fixed, this trading is limited because of the transaction cost of bartering with other dealers and thin markets resulting from the large variety of cars.

² See Section 2 for an explanation of why substitute inventory can be considered exogenous in the short to medium run.

³ The standard setup in which the price–inventory relationship has been studied is a situation in which prices are set by a monopolist who has to sell a given stock by a deadline. In that situation, Gallego and van Ryzin (1994) show that the optimal price is nonincreasing in the remaining inventory and nondecreasing in the remaining time until the deadline. The situation we describe, in which a dealer can reorder inventory but there is a lag between ordering and resupply, is a straightforward extension to the standard setup. In the online appendix we provide an example of such a model.

⁴ Note that, although we expect the slope to be smaller in more competitive markets, we still expect more competitive markets to have lower prices overall.

⁵ However, they can exchange vehicles with other dealers. In the empirical analysis, we control for interdealer trades. See Section 2.3 for a discussion of dealer trades.

⁶ Because of our focus on the dealer's short-run pricing problem, we do not address the interesting issue raised in Carlton (1978) and Dana (2001), namely that a firm chooses both a price at which to sell its good and a level of availability. In the context of car dealers, this involves the dealer choosing to have a full or limited selection on the lot and then compensating customers for the benefit or cost of that choice with the price of the car. Empirically, because all the estimations in our paper include dealer fixed effects, we are effectively controlling for the strategic choice of availability on the part of the dealer by estimating the effect of inventory off intradealer inventory levels.

⁷ A particular car that is scheduled to be delivered can be reserved with a down payment, which functions as a contract promising a future sale at a specific price. This down payment is often relatively small, so the customer still has considerable freedom to choose another car. According to an industry source, Americans do not employ this strategy as much as Europeans: fewer than 3% of Americans preorder a car, whereas, in some countries in Europe, as many as 50% of consumers do.

⁸ We find that the price–inventory effect does not differ significantly by dealer sales volume (not reported).

⁹ One might ask why substitute inventory should be thought of measuring market power instead of bargaining power. The bargaining literature distinguishes between the parties' outside options and their bargaining power. Whether a party has outside options reflects whether the party has positive disagreement payoffs. Bargaining power depends on the parties' ability to commit to their offer. The two concepts (commitment and outside options) are quite different. Bargaining power is usually thought of as the relative patience of the parties, whereas outside options reflect whether there are alternative buyers or sellers. In our setting, the availability of substitute inventory directly affects the outside option of consumers and, therefore, the market power of a focal dealer. Instead, a dealer's bargaining power is best thought of as the dealer's patience relative to the consumer in arriving at a deal.

¹⁰ Our data contain 141 such markets across the United States.

¹¹ In our data, 56% of transactions come from consumers who reside within 10 miles, 80% within 20 miles, 88% within 30 miles, and 92% within 40 miles of the dealership from which they buy the car.

¹² We show in Section 4.3.2 that our results are robust to classifying substitute inventory with more granularity.

¹³ In multiple interviews, we asked repeatedly whether there were any exceptions to basing transfer payments on invoice prices. No interviewee had heard of any other practice.

¹⁴ Dealer-installed accessories that contribute to the resale value include items such as upgraded tires or a sound system but exclude options such as undercoating or waxing.

¹⁵ Holdback is an amount the manufacturer adds to the vehicle invoice that is later refunded to the dealer, typically 2%–5% of the invoice price.

¹⁶ This is the finest car description available in our data. Notice that we measure inventory at a slightly more aggregate level by combining different engine sizes, trim levels, and transmission type.

¹⁷ For robustness, we ran the baseline specifications for the entire sample and obtained coefficients and p -values that were similar (unreported). None of our conclusions change.

¹⁸ For a monopolist, the standard intuition is that, as the date nears, the monopolist's opportunity cost of selling the remaining cars on the lot falls, holding constant the level of inventory, because soon the dealer will be restocked. Thus, as the number of days to

resupply drops, the dealer should be more willing to discount the car to a customer with a low valuation.

¹⁹ For a thorough analysis of the effects of demographics on car prices, see Scott Morton et al. (2003)

²⁰ Of course, a competent dealer in this situation would try to adjust inventory in the long run, and so this story really only applies if this proves difficult or if the shock is transitory.

²¹ For robustness, we also define market power by using a 20-mile radius around a focal dealer. In 80% of the transactions, consumers reside within 20 miles of the dealership from which they buy a car. The results are consistent with the results for the 30-mile radius but are unreported in the interest of space.

²² Note that the figure only graphs the interaction effects, not the main effects of substitute inventory.

²³ This is noted in <https://www.prnewswire.com/news-releases/study-americans-feel-taken-advantage-of-at-the-car-dealership-300301866.html>, accessed December 6, 2020.

²⁴ For example, <https://www.consumer.ftc.gov/articles/0209-buying-new-car>, accessed December 6, 2020.

²⁵ For example, "U.S. automakers suffer from dealer networks that are too big and bogged down by weak performers, said Roger Penske, the racing legend and billionaire businessman who heads one of the largest chains of auto dealerships, reports the Associated Press. Some of the competitors come in and will have less dealers that have larger scale, who then have the ability to spend more money in the marketplace," Penske said." This can be found at <https://leftlanenews.com/2006/01/04/do-the-big-three-have-too-many-dealers/>, accessed December 6, 2020.

References

- Amihud Y, Mendelson H (1989) Inventory behaviour and market power: An empirical investigation. *Internat. J. Indust. Organ.* 7(2): 269–280.
- Anderson SP, Schneider Y (2007) Price competition with revenue management and costly consumer search. Working paper, University of Virginia, Charlottesville.
- Borenstein S, Rose N (1994) Competition and price dispersion in the U.S. airline industry. *J. Political Econom.* 102(4):653–683.
- Busse MR, Silva-Risso J, Zettelmeyer F (2006) \$1000 cash back: The pass-through of auto manufacturer promotions. *Amer. Econom. Rev.* 96(4):1253–1270.
- Cachon GP, Olivares M (2010) Drivers of finished-goods inventory in the U.S. automobile industry. *Management Sci.* 56(1):202–216.
- Carlton D (1978) Market behavior with demand uncertainty and price inflexibility. *Amer. Econom. Rev.* 68(4):571–587.
- Chen N (2018) Perishable good dynamic pricing under competition: An empirical study in the airline markets. Discussion paper, The National University of Singapore, School of Computing.
- Chen X, Simchi-Levi D (2012) Pricing and inventory management. Özer Ö, Phillips R, eds. *The Oxford Handbook of Pricing Management* (Oxford University Press, Oxford, UK), 784–824.
- Copeland A, Hall G (2011) The response of prices, sales, and output to temporary changes in demand. *J. Appl. Econometrics* 26(2): 232–269.
- Copeland A, Dunn W, Hall G (2011) Inventories and the automobile market. *RAND J. Econom.* 42(1):121–149.
- Dana J (2001) Competition in price and availability when availability is unobservable. *RAND J. Econom.* 32(4):497–513.
- Dana JR, Williams KR (2020) Intertemporal price discrimination in sequential quantity-price games. Cowles Foundation Discussion Paper No. 2136R3, Northeastern and Yale University. <https://cowles.yale.edu/sites/default/files/files/pub/d21/d2136-r3.pdf>.

- Dudey M (1992) Dynamic Edgeworth-Bertrand competition. *Quart. J. Econom.* 107(4):1461–1477.
- Gallego G, Hu M (2014) Dynamic pricing of perishable assets under competition. *Management Sci.* 60(5):1241–1299.
- Gallego G, van Ryzin G (1994) Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* 40(8):999–1020.
- Graddy K, Hall G (2011) A dynamic model of price discrimination and inventory management at the Fulton fish market. *J. Econom. Behav. Organ.* 80(1):6–19.
- Hall G, Rust J (2000) An empirical model of inventory investment by durable commodity intermediaries. *Carnegie-Rochester Conf. Ser. Public Policy* 52:171–214.
- Levin Y, McGill J, Nediak M (2009) Dynamic pricing in the presence of strategic consumers and oligopolistic competition. *Management Sci.* 55(1):32–46.
- Lin KY, Sibdari SY (2009) Dynamic price competition with discrete customer choices. *Eur. J. Oper. Res.* 197(3):969–980.
- Liu Q, Zhang D (2013) Dynamic pricing competition with strategic customers under vertical product differentiation. *Management Sci.* 59(1):84–101.
- Mantin B, Granot D, Granot F (2011) Dynamic pricing under first order Markovian competition. *Naval Res. Logist.* 58(6): 608–617.
- Martínez-de Albéniz V, Talluri K (2011) Dynamic price competition with fixed capacities. *Management Sci.* 57(6):1078–1093.
- Mookherjee R, Friesz TL (2008) Pricing, allocation, and overbooking in dynamic service network competition when demand is uncertain. *Production Oper. Management* 17(4):455–474.
- Olivares M, Cachon GP (2009) Competing retailers and inventory: An empirical investigation of General Motors' dealerships in isolated U.S. markets. *Management Sci.* 55(9):1586–1604.
- Scott Morton F, Zettelmeyer F, Silva-Risso J (2003) Consumer information and discrimination: Does the internet affect the pricing of new cars to women and minorities? *Quant. Marketing Econom.* 1(1):65–92.
- Siegart C, Ulbricht R (2020) Dynamic oligopoly pricing: Evidence from the airline industry. *Internat. J. Indust. Organ.* 71:102639.
- Xu X, Hopp WJ (2006) A monopolistic and oligopolistic stochastic flow revenue management model. *Oper. Res.* 54(6): 1089–1109.