

How Market Power Affects Dynamic Pricing:
Evidence from Inventory Fluctuations at Car Dealerships^{*}

Ayelet Israeli

Harvard University

Fiona Scott Morton

Yale University and NBER

Jorge Silva-Risso

University of California, Riverside

Florian Zettelmeyer

Northwestern University and NBER

December 2020

^{*} We are grateful for helpful comments from Eva Ascarza, Severin Borenstein, Dennis Carlton, Jose Silva, K. Sudhir, Candi Yano, Jan Van Miegen, Martin Lariviere, participants at the Junior Faculty in Marketing Science Forum, the NBER IO Program Meeting, the Marketing in Israel Conference, and seminar participants at Harvard Business School, IDC Herzliya, Northwestern University, Stanford University, UC Berkeley, the University of Chicago, the University of East Anglia, and the University of Rochester. We are especially grateful to Meghan Busse for extensive comments and to Thomas Hubbard for the suggestion to look into inventories. E-mail: aisraeli@hbs.edu, fiona.scottmorton@yale.edu, jorge.silva-risso@ucr.edu, f-zettelmeyer@kellogg.northwestern.edu

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

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, i.e., 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.

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 their 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, i.e., the quantity of substitute inventory in their 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,

¹Even if inter-dealer 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 due to the large variety of cars.

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 a sales effect and a service effect 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 due to price discrimination.

Although these papers have provided convincing evidence on the way that market power affects inventories or price, 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 this 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 inter-temporal 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 re-order 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 her expectation or

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

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

“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 where there is *no* substitute inventory—the situation a monopolist dealer faces. As described above, lower inventory should lead to higher prices. Now consider the other extreme case where there is ample *perfectly substitutable* inventory. The dealer will not be 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 due to, for example, a

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

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 (Mantin, Granot, and Granot 2011, Mookherjee and Friesz 2008, Xu and Hopp 2006, Anderson and Schneider 2007, Dudey 1992, Martínez-de Albéniz and Talluri 2011). Lin and Sibdari (2009) show that under competition, the optimal price for a product need not be non-decreasing 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, McGill, and Nediak (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 his customers and displays substantial fluctuation in day-to-day inventory of different products. Copeland, Dunn, and Hall (2011) model the optimal pricing and production decisions of auto manufacturers which 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

Williams (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 U.S. 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, re-ordering 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 she wants on the lot, she 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. While 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

⁵However, they can exchange vehicles with other dealers. In the empirical analysis we control for inter-dealer trades. Please see section 2.3 for a discussion of dealer trades.

⁶Because 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 would involve the dealer choosing to have a full or limited selection on her 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 intra-dealer inventory levels.

⁷A 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 pre-order a car, whereas in some countries in Europe as many as 50% of consumers do.

pricing mechanism in the web 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 will 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 will be 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 will be 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 which 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 dataset comprises car purchases for 16 years from January 1, 1998 to December 31, 2013. Figure 4 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 which inventory should take in determining prices. One might expect that inventory will 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 1 for small dealerships (see the top panel of Figure 5 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 which 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, 15, and one for 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 5 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

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

products A and B. Product A is sold in a competitive market while the firm is a monopolist in the market for product B. We would 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 where 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 inter-temporal 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 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 in 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 7 presents a dealer’s own inventory levels compared to the local substitute inventory levels (based on the DMA measure) over time for the same Honda dealer from Figure 4. As the figure shows, local substitute inventory levels vary widely over time within car-dealer combination, effectively leading to variation in market power for a dealer selling a focal car. To demonstrate

⁹One 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.

¹⁰Our data contain 141 such markets across the US.

¹¹In 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.

the variation in market power across car-dealer combinations, for each combination we compute the standard deviation of local substitute inventory. Figure 8 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 where 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 9 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 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 dealer, and also allows for a large variation in market power across cars within 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 since 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. 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

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

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 which 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 which 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 mileage” or “high mileage.” We construct a mileage cut-off 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 will contain traded cars. We then define a *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 above. Because every requested trade results in a received trade at the reciprocating dealer, we define a car as a *TradeReceived* if it had an odometer reading that exceeded the same 95th quantile, was not a *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 under 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 dataset and for the restricted dataset 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.

¹³In 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.

We will 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 less before the introduction of the next model year. We omit these transactions from the dataset 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 dataset is the price that the customer pays for the vehicle, including factory-installed accessories and options and the 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 will control for many of the factors that contribute to the price of a car, it will not control for the factory- and dealer-installed options which 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

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

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

¹⁶This 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.

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 dataset. This percentage deviation, called *VehicleCost*, will be 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 above). 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 1 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 will pick up their effect on prices with these variables.

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, while 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 dataset. By excluding car types that sold fewer than 2991 times over the model year nationwide, and dealers who sold fewer than 615 cars over the sample period, we reduced 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 over USD 100,000. Our final dataset contains 4,903,122 observations. Summary statistics for the dataset 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 due to 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 due to 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 will 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 will be relying neither on variation across dealerships, nor variation across cars, nor variation across 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 will 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 will discuss both approaches in more detail in the next sections.

¹⁷For 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.

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:

$$\ln(\text{Price}_i) = X_i\alpha + D_i\beta + I_i\gamma + \epsilon_i \quad (1)$$

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 her census block group. To this basic specification we add a matrix I which contains various inventory-related explanatory variables such as measures of inventory, days to resupply, 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, 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 car 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 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 non-decreasing in time-to-go.¹⁸

Highlighting some other results, we find that consumers pay a lower price (0.18%) for a vehicle which was requested from another dealership (*TradeCar*). 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 due to a potential endogeneity of prices and inventory levels due to demand shocks. We use a sequence of fixed effects

¹⁸For a monopolist, the standard intuition is that, as the date nears, the monopolist’s opportunity cost of selling the remaining cars on her 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.

¹⁹For a thorough analysis of the effects of demographics on car prices, please see Scott Morton, Zettelmeyer, and Silva-Risso (2003)

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 due to 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 above 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 will lead 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: while 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 \times 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

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

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 will absorb 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 which 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 0.029% (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 their 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, i.e., the quantity of substitute inventory in their 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 in 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:

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

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 her 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_i , directly. Finally, the I_i matrix contains the inventory-related controls such as days to resupply, and a trade-requested indicator.

²¹For 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.

We test two predictions using this equation. First, 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 will be 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) are 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.

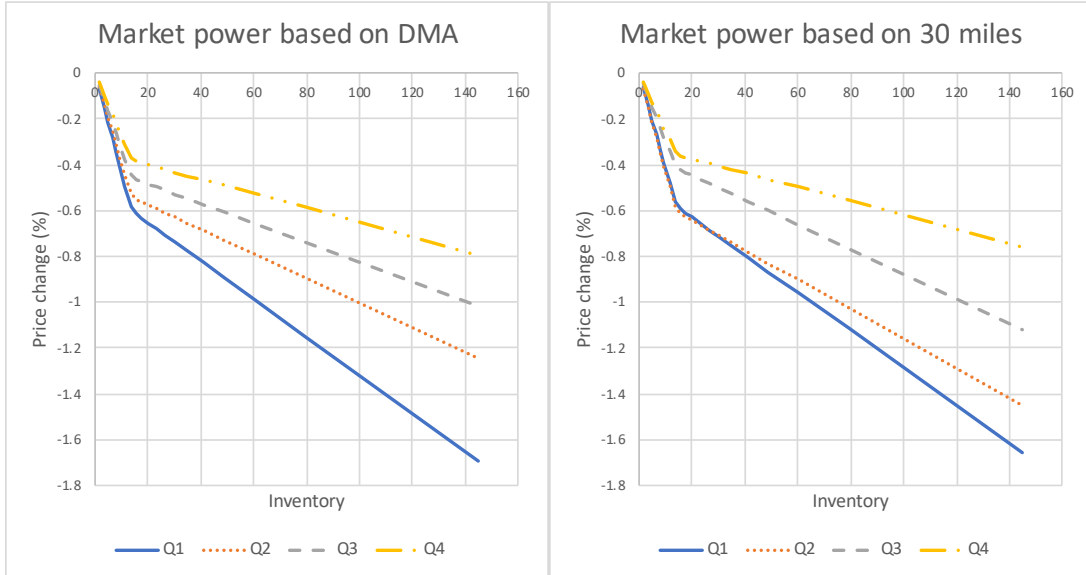
Our second and key prediction is tested using the interaction between the inventory spline and the market power quartiles. We find that more substitute inventory in the DMA leads to a weaker (less negative) price–inventory relationship. The interaction coefficients are statistically different from each other (except for the interactions of local q3 and local 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 1 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,

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

Figure 1: Market power price–inventory relationship slopes



the more inventory competing retailers have of the same type. These results are robust to the definition of local selling area as 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 on the trade-in, dealers and sales people earn a margin from car financing and insurance (“F&I margins”). In this subsection, we test whether F&I margins, another component of price, are also affected by inventories. During a new car sale, after the customer agreed 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—will offer 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 while 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 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. Since 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 above-median level of the focal dealer’s inventory. For below-median focal inventory, the hypothesized interaction is not present. Figure 2 illustrates these results.

Figure 2: Robustness: more granular definition of a car



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 due to above-median focal inventory but not below-median inventory. One interpretation is that consumers are willing to substitute to very similar cars (which in our narrower inventory definition are defined as a different car) when inventories are low.

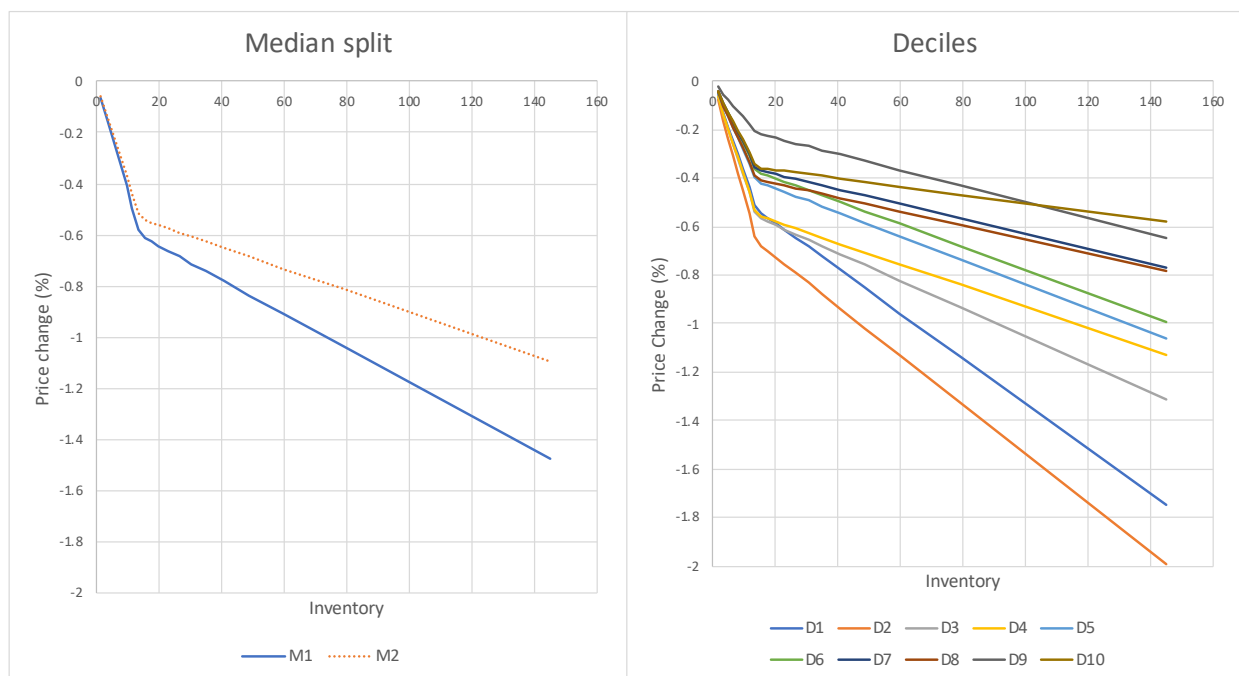
4.3.2 Do our results depend on the granularity of our market power measure?

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 splits and decile splits of the substitute inventory. For clarity of presentation, we present the results in a figure.

Figure 3 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 where substitute inventory

is exactly 0, i.e., the focal dealer is a monopoly in their local market, and d2 includes cases where substitute inventory contains 1–9 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.

Figure 3: Robustness: different definitions of market power



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 subsection 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 “offline information” periods based on whether or not 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 column 2 and column 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 “offline” 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, Silva-Risso, and Zettelmeyer (2006) for the relevant literature). In these models reducing the information asymmetry of a party will allow that party to obtain a larger share of the surplus in the negotiation. In our setting, consumers are initially uninformed about inventories, and they 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 to those in the “offline” 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 investigate how market power affects firms' ability to dynamically price. To do do, 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) are 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 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, consumer strategies, 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—will like vary 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 from 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 will 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 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

²³<https://www.prnewswire.com/news-releases/study-americans-feel-taken-advantage-of-at-the-car-dealership-300301866.html>, accessed on 12/6/2020.

²⁴For example, <https://www.consumer.ftc.gov/articles/0209-buying-new-car>, accessed on 12/6/2020.

²⁵For 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. “<https://leftlanenews.com/2006/01/04/do-the-big-three-have-too-many-dealers/> accessed on 12/6/2020.

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.

References

- AMIHUD, Y., AND H. MENDELSON (1989): “Inventory behaviour and market power: An empirical investigation,” *International Journal of Industrial Organization*, 7, 269–280.
- ANDERSON, S. P., AND Y. SCHNEIDER (2007): “Price competition with revenue management and costly consumer search,” Working paper, University of Virginia.
- AUSUBEL, L. M., AND R. J. DENECKERE (1998): “Bargaining and Forward Induction,” Unpublished Paper, Department of Economics, University of Maryland.
- BORENSTEIN, S., AND N. ROSE (1994): “Competition and Price Dispersion in the U.S. Airline Industry,” *Journal of Political Economy*, 102(4), 653–683.
- BRESNAHAN, T. F., AND V. A. RAMEY (1994): “Output Fluctuations at the Plant Level,” *Quarterly Journal of Economics*, 109(3), 593–624.
- BUSSE, M. R., J. SILVA-RISSO, AND F. ZETTELMEYER (2006): “\$1000 Cash Back: The Pass-Through of Auto Manufacturer Promotions,” *American Economic Review*, 96(4), 1253–1270.
- CACHON, G. P., AND M. OLIVARES (2010): “Drivers of Finished-Goods Inventory in the U.S. Automobile Industry,” *Management Science*, 56, 202–216.
- CARLTON, D. (1978): “Market Behavior with Demand Uncertainty and Price Inflexibility,” *American Economic Review*, 68, 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., AND D. SIMCHI-LEVI (2012): “Pricing and inventory management,” in *The Oxford handbook of pricing management*, ed. by Ö. Özer, and R. Phillips, pp. 784–824. Oxford University Press.
- COPELAND, A., W. DUNN, AND G. HALL (2011): “Inventories and the Automobile Market,” *RAND Journal of Economics*, 42, 121–149.
- COPELAND, A., AND G. HALL (2011): “The Response of Prices, Sales, and Output to Temporary Changes in Demand,” *Journal of Applied Econometrics*, 26, 232–269.
- DANA, J. (2001): “Competition in price and availability when availability is unobservable,” *RAND Journal of Economics*, 32(4), 497–513.
- DANA, J. R., AND K. R. WILLIAMS (2020): “Intertemporal Price Discrimination in Sequential Quantity-Price Games,” Discussion paper, Northeastern and Yale University.
- DUDEY, M. (1992): “Dynamic Edgeworth-Bertrand Competition,” *The Quarterly Journal of Economics*, 107, 1461–1477.

- GALLEGO, G., AND M. HU (2014): “Dynamic Pricing of Perishable Assets Under Competition,” *Management Science*, 60(5), 1241–1299.
- GALLEGO, G., AND G. V. RYZIN (1994): “Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons,” *Management Science*, 40(8), 999–1020.
- GRADDY, K., AND G. HALL (2011): “A Dynamic Model of Price Discrimination and Inventory Management at the Fulton Fish Market,” *Journal of Economic Behavior and Organization*, 80, 6–19.
- HALL, G., AND J. RUST (2000): “An Empirical Model of Inventory Investment by Durable Commodity Intermediaries,” *Carnegie-Rochester Conference Series on Public Policy*.
- KENNAN, J., AND R. WILSON (1993): “Bargaining with Private Information,” *Journal of Economic Literature*, 31(1), 45–104.
- KUO, C.-W., H.-S. AHN, AND G. AYDIN (2011): “Dynamic Pricing of Limited Inventories When Customers Negotiate,” *Operations Research*, 59, 882–897.
- LEVIN, Y., J. MCGILL, AND M. NEDIAK (2009): “Dynamic pricing in the Presence of Strategic Consumers and Oligopolistic Competition,” *Management Science*, 55(1), 32–46.
- LIN, K. Y., AND S. Y. SIBDARI (2009): “Dynamic price competition with discrete customer choices,” *European Journal of Operations Research*, 197, 969–980.
- LIU, Q., AND D. ZHANG (2013): “Dynamic Pricing Competition with Strategic Customers Under Vertical Product Differentiation,” *Management Science*, 59, 84–1010.
- MANTIN, B., D. GRANOT, AND F. GRANOT (2011): “Dynamic pricing under first order Markovian competition,” *Naval Research Logistics*, 58, 608–617.
- MARTÍNEZ-DE ALBÉNIZ, V., AND K. TALLURI (2011): “Dynamic Price Competition with Fixed Capacities,” *The Quarterly Journal of Economics*, 57, 1078–1093.
- MOOKHERJEE, R., AND T. L. FRIESZ (2008): “Pricing, Allocation, and Overbooking in Dynamic Service Network Competition When Demand Is Uncertain,” *Production and Operation Management*, 17, 455–474.
- OLIVARES, M., AND G. P. CACHON (2009): “Competing Retailers and Inventory: An Empirical Investigation of General Motors’ Dealerships in Isolated U.S. Markets,” *Management Science*, 55, 1586–1604.
- PAPASTAVROU, J. D., S. RAJAGOPALAN, AND A. J. KLEYWEGT (1996): “The Dynamic and Stochastic Knapsack Problem with Deadlines,” *Management Science*, 42(12), 1706–1718.
- SCOTT MORTON, F., F. ZETTELMEYER, AND J. SILVA-RISSO (2003): “Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities?,” *Quantitative Marketing and Economics*, 1(1), 65–92.

- SIEGART, C., AND R. ULBRICHT (2020): “Dynamic Oligopoly Pricing: Evidence from the Airline Industry,” *International Journal of Industrial Organization*, 71.
- XU, X., AND W. J. HOPP (2006): “A Monopolistic and Oligopolistic Stochastic Flow Revenue Management Model,” *Operations Research*, 54, 1089–1109.
- YANO, C. A., AND S. M. GILBERT (2003): “Coordinated Pricing and Production/Procurement Decisions: A Review,” in *Managing Business Interfaces: Marketing, Engineering and Manufacturing Perspectives*, ed. by A. Chakravarty, and J. Eliashberg. Kluwer Academic Publishers.

Table 1: Summary statistics

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

Table 2: Basic Result: Price effects of inventory[†]

	Fixed Effects			
	Car, Dealer, Month	Car×Dealer Month	Car×Dealer, Month×Segment×DMA	Car, Dealer, Week×Subsegment×DMA
Inventory (1–14)	-.039** (.00087)	-.046** (.0011)	-.038** (.0012)	-.027** (.001)
Inventory (15+)	-.0057** (.00013)	-.005** (.00022)	-.0033** (.00024)	-.0037** (.00016)
DaysToResupply	-.0022** (.00023)	-.00097** (.0003)	-.0001 (.0003)	-.0009 (.00026)
TradedCar	-.18** (.0087)	-.22** (.0092)	-.23** (.0091)	-.19** (.009)
Tradein	2.6** (.006)	2.6** (.0061)	2.5** (.0061)	2.6** (.0062)
VehicleCost	84** (.092)	86** (.099)	86** (.098)	85** (.093)
Model Age 5–13 Months	.047** (.011)	.045** (.011)	-.061** (.013)	.013 (.015)
Model Age > 14 Months	-.043* (.022)	-.028 (.022)	-.089** (.025)	-.014 (.029)
Weekend	.074** (.0062)	.085** (.0063)	.082** (.0062)	.058** (.0065)
EndOfMonth	-.52** (.0068)	-.5** (.0068)	-.49** (.0068)	-.17** (.015)
EndOfYear	-.19** (.02)	-.16** (.02)	-.16** (.02)	-.098* (.04)
Observations	4,903,122	4,903,122	4,903,122	4,903,122
Adj. R-squared	0.956	0.959	0.960	0.958

[†] Significant at 10%; [*] significant at 5%; ** significant at 1%. Robust standard errors in parentheses.

[†] Unreported are the constant term, car, dealer, and month fixed effects (column 1); car*dealer and monthly fixed effect (column 2); car, dealer, and month*segment*DMA fixed effect (column 3); and car, dealer, week*subsegment*DMA fixed effect (column 4); the demographic variables reported in Table 1. .

All coefficients are multiplied by 100.

Table 3: Main Results: Local and focal inventory effects on prices[†]

Dep. Var. ln(price)	(1) Local inventory defined at DMA level	(2) Local inventory defined at 30 miles radius
Local q2	-.27** (.024)	-.19** (.024)
Local q3	-.64** (.031)	-.6** (.031)
Local q4	-1** (.045)	-.98** (.045)
Local q1×Inventory (1–14)	-.043** (.0018)	-.041** (.0018)
Local q2×Inventory (1–14)	-.039** (.0018)	-.043** (.0018)
Local q3×Inventory (1–14)	-.033** (.0021)	-.029** (.0021)
Local q4×Inventory (1–14)	-.027** (.0029)	-.025** (.003)
Local q1×Inventory (15+)	-.0084** (.00049)	-.0082** (.00053)
Local q2×Inventory (15+)	-.0054** (.00051)	-.0065** (.00052)
Local q3×Inventory (15+)	-.0043** (.00041)	-.0054** (.00039)
Local q4×Inventory (15+)	-.0031** (.00026)	-.0031** (.00026)
DaysToResupply	-.0009** (.0003)	-.00083** (.0003)
TradedCar	-.22** (.0092)	-.22** (.0092)
Tradein	2.6** (.0061)	2.6** (.0061)
VehicleCost	86** (.099)	86** (.099)
Model Age 5–13 Months	.069** (.011)	.066** (.011)
Model Age > 14 Months	-.019 (.022)	-.02 (.022)
Weekend	.085** (.0063)	.088** (.0063)
EndOfMonth	-.49** (.0068)	-.49** (.0068)
EndOfYear	-.17** (.02)	-.17** (.02)
Observations	4,903,122	4,903,122
Adj. R-squared	0.959	0.959

⁺ Significant at 10%; [*] significant at 5%; ** significant at 1%. Robust standard errors in parentheses.

[†] Unreported are the constant term, car*dealer, monthly fixed effect, and the demographic variables reported in Table 1.

All coefficients are multiplied by 100.

Table 4: Financing and insurance results[†]

Dep. Var. ln(F&I)	(1)	(2)
Inventory (1–14)	-.005** (.00048)	
Inventory (15+)	-.0003** (.00008)	
DaysToResupply	-.00003 (.0001)	-.00003 (.0001)
local q2		-.03** (.01)
local q3		-.07** (.013)
local q4		-.095** (.02)
local q1×Inventory (1–14)		-.007** (.00072)
local q2×Inventory (1–14)		-.004** (.00074)
local q3×Inventory (1–14)		-.0025** (.00089)
local q4×Inventory (1–14)		-.0022 (.0014)
local q1×Inventory (15+)		-.0001 (.00017)
local q2×Inventory (15+)		-.00097** (.00018)
local q3×Inventory (15+)		-.00055** (.00015)
local q4×Inventory (15+)		-.00018+ (.0001)
TradedCar	-.11** (.0043)	-.11** (.0043)
Tradein	.098** (.0024)	.098** (.0024)
VehicleCost	.26** (.025)	.26** (.025)
Model Age 5–13 Months	-.029** (.0047)	-.027** (.0047)
Model Age > 14 Months	-.033** (.0089)	-.032** (.0089)
Weekend	.064** (.0025)	.064** (.0025)
EndOfMonth	-.058** (.0028)	-.057** (.0028)
EndOfYear	.0038 (.0082)	.0037 (.0082)
Observations	2,758,335	2,758,335
Adj. R-squared	0.211	0.212

⁺ Significant at 10%; [*] significant at 5%; ** significant at 1%. Robust standard errors in parentheses.

[†] Unreported are the constant term, car*dealer, monthly fixed effect, and the demographic variables reported in Table 1.

All coefficients are multiplied by 100.

Table 5: Inventory definition and inventory information[†]

Dep. Var. ln(price)	(1) Narrower Inventory Def.	(2) Online Information	(3) Offline Information
local q2	-.18** (.021)	-.35** (.028)	-.13** (.043)
local q3	-.43** (.028)	-.82** (.037)	-.23** (.06)
local q4	-.76** (.044)	-1.2** (.051)	-.66** (.093)
local q1×Inventory (1–14)	-.054** (.0036)	-.049** (.0021)	-.033** (.0036)
local q2×Inventory (1–14)	-.065** (.0033)	-.039** (.0021)	-.04** (.0033)
local q3×Inventory (1–14)	-.065** (.0037)	-.03** (.0025)	-.041** (.0041)
local q4×Inventory (1–14)	-.07** (.006)	-.025** (.0034)	-.033** (.0063)
local q1×Inventory (15+)	-.011** (.00086)	-.009** (.00054)	-.0056** (.0013)
local q2×Inventory (15+)	-.0074** (.00097)	-.0067** (.0006)	-.0045** (.001)
local q3×Inventory (15+)	-.0062** (.0007)	-.0044** (.00047)	-.007** (.00085)
local q4×Inventory (15+)	-.0034** (.00041)	-.0036** (.0003)	-.0033** (.0006)
DaysToResupply	-.00025 (.00016)	-.00057+ (.00034)	-.00055 (.00065)
TradedCar	-.23** (.0092)	-.21** (.011)	-.25** (.017)
Tradein	2.6** (.0061)	2.2** (.0071)	3.5** (.012)
VehicleCost	86** (.099)	87** (.12)	84** (.19)
Model Age 5–13 Months	.058** (.011)	.042** (.013)	.065** (.023)
Model Age > 14 Months	-.021 (.022)	-.017 (.026)	-.056 (.043)
Weekend	.081** (.0063)	.088** (.0072)	.081** (.012)
EndOfMonth	-.5** (.0068)	-.54** (.008)	-.36** (.013)
EndOfYear	-.17** (.02)	-.19** (.024)	-.15** (.037)
Observations	4,903,122	3,584,401	1,381,721
Adj. R-squared	0.959	0.958	0.960

⁺ Significant at 10%; [*] significant at 5%; ** significant at 1%. Robust standard errors in parentheses.

[†] Unreported are the constant term, car*dealer, monthly fixed effect, 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.

Figure 4: Example of inventory movement for three cars at one Honda dealer in Jan–Feb 2013

Dealer 24749 (Honda)

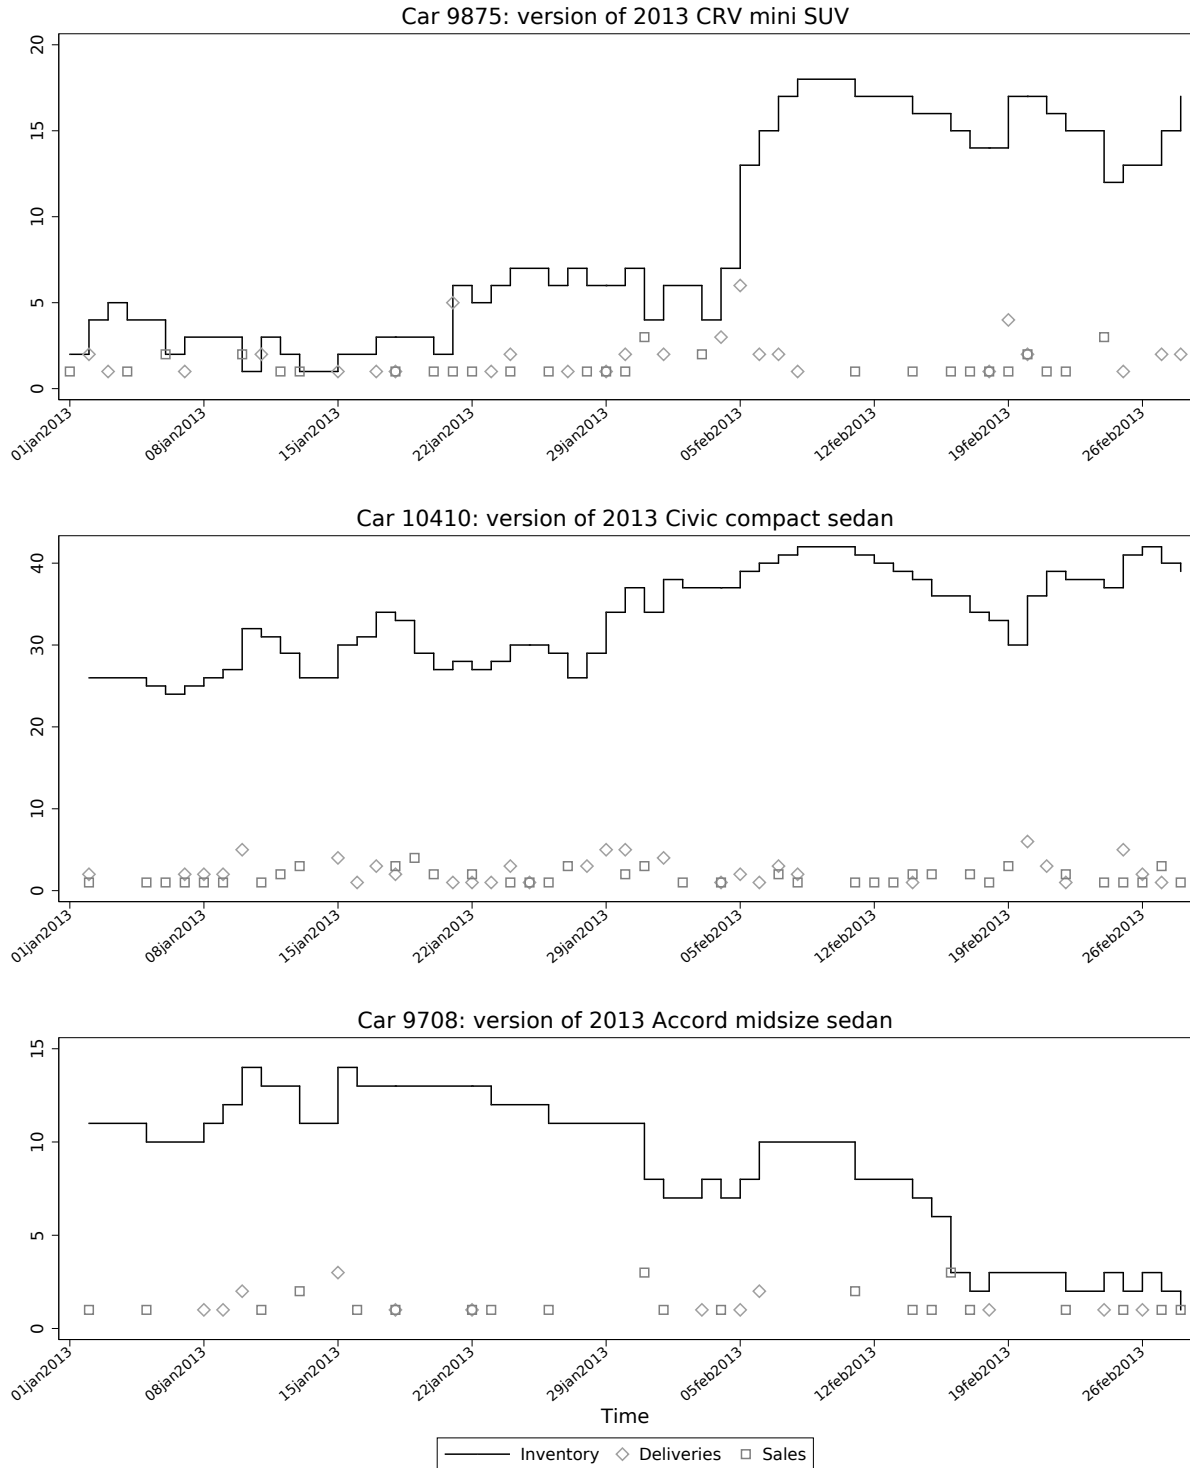


Figure 5: Distribution of daily inventories (at the inventory “car” level)

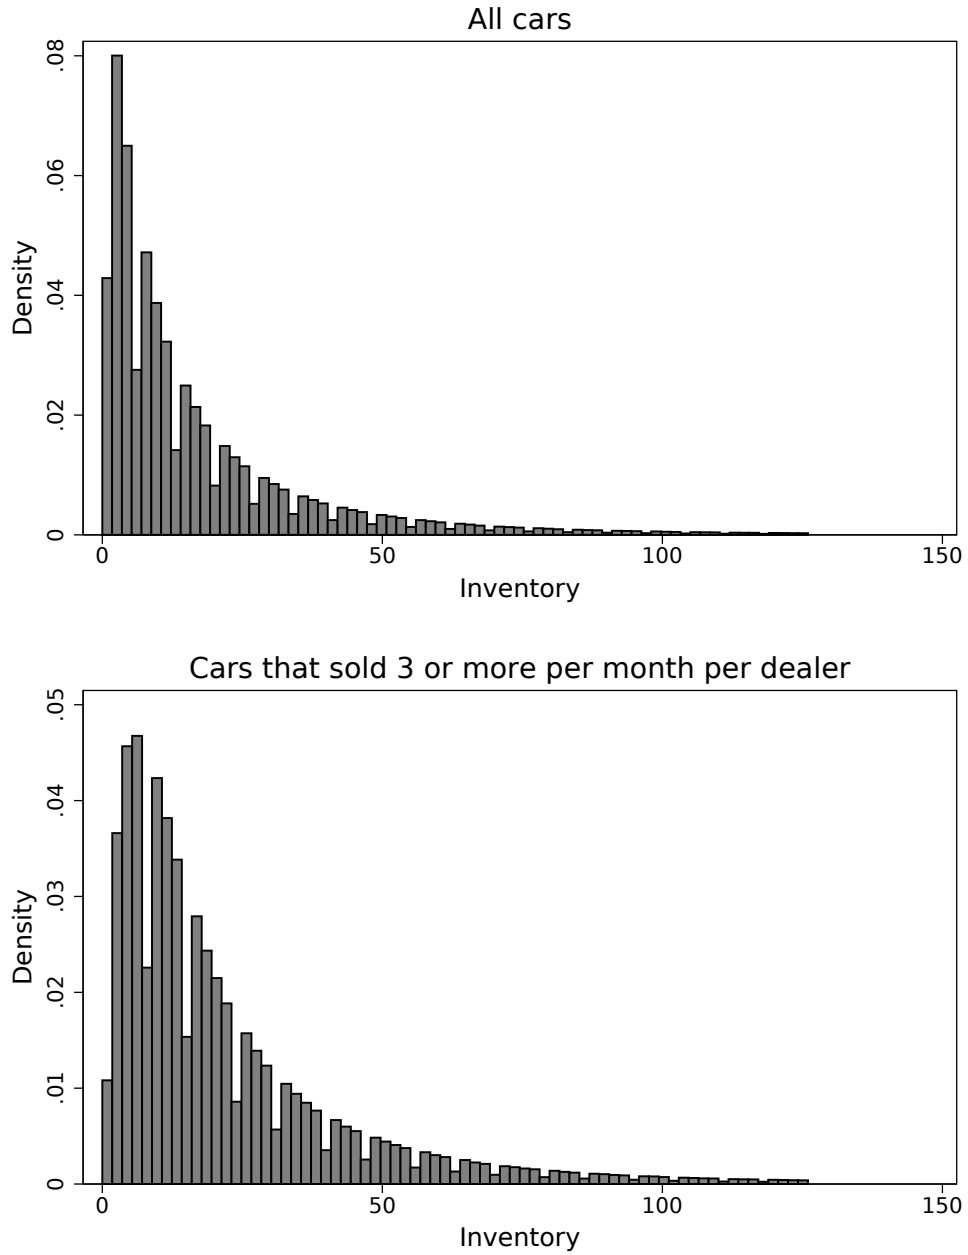


Figure 6: Distribution of daily “days to resupply” (at the inventory “car” level)

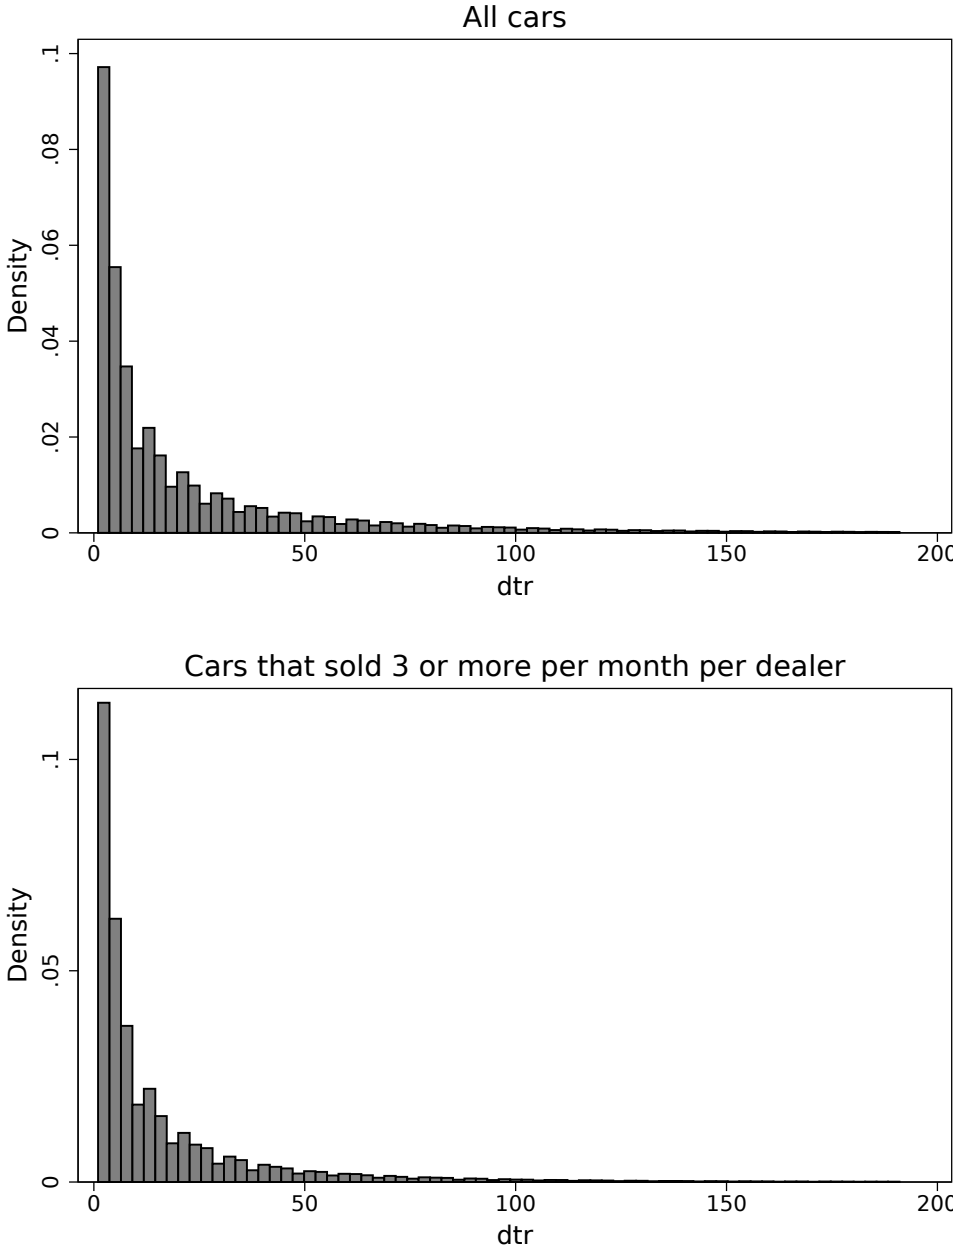


Figure 7: Example of focal and local market inventory movement for three cars at one Honda dealer in Jan–Feb 2013

Dealer 24749 (Honda)

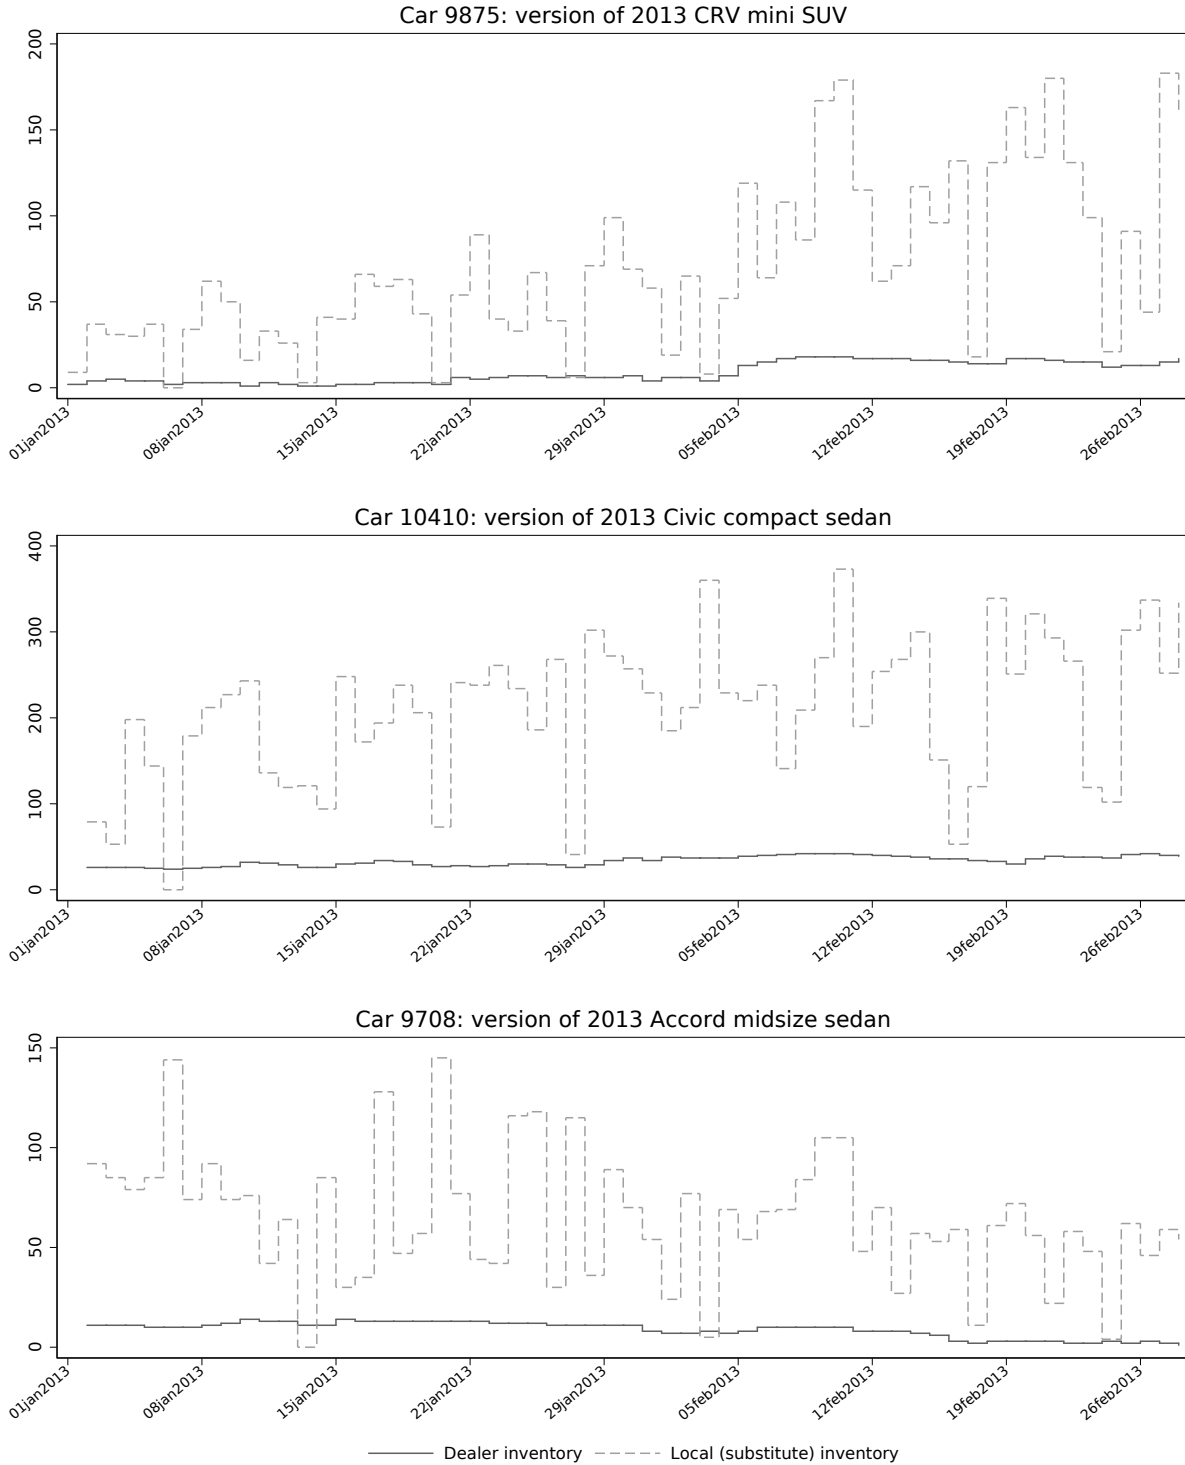


Figure 8: Distribution of standard deviations of local substitute inventories (at the dealership-car level)

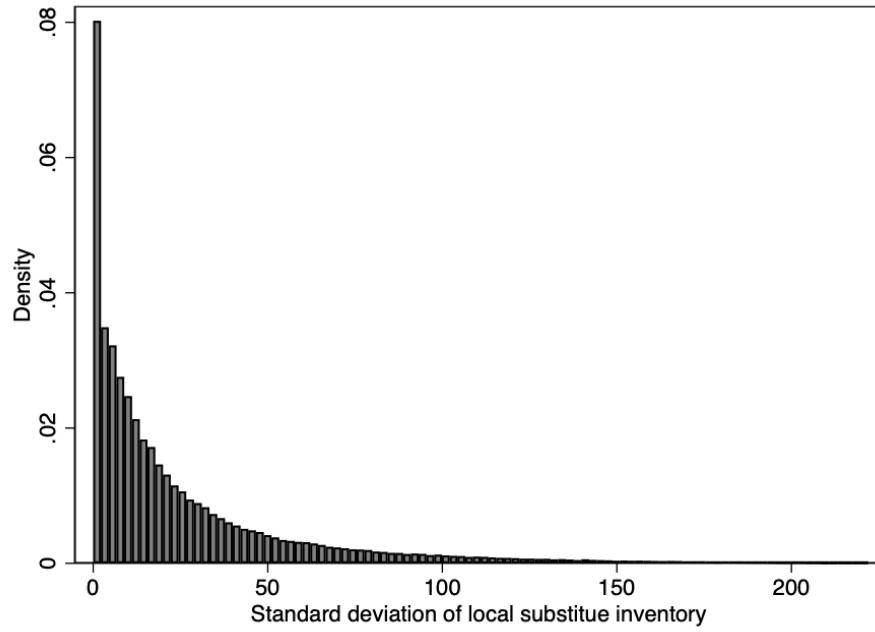
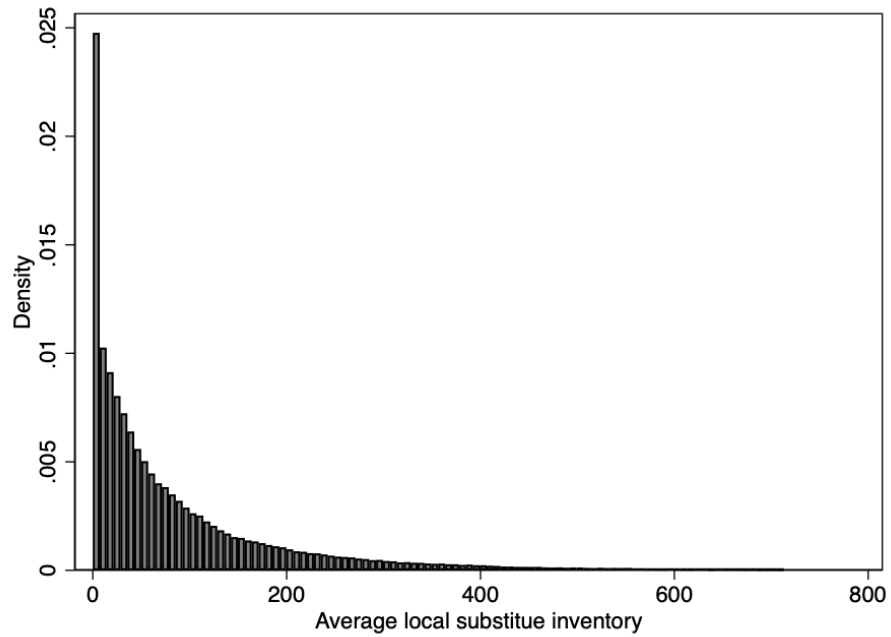


Figure 9: Distribution of average local substitute inventories (at the dealership-car level)



WEB APPENDIX

1 Model of Inventory and prices

To develop an intuition for the relation of prices and inventory, we set up a simple infinite-horizon model of dealer pricing with stochastic demand. We then derive dealer pricing as a function of inventory for a specific example. We use the insights from the model to derive empirical predictions.

We do not attempt to discuss or solve for the optimal method to allocate cars to dealers. The analysis in the paper is ultimately identified from small daily movements in a dealer’s inventory while controlling for the characteristics of the car, demand conditions, buyer characteristics, and the arrival time of any new inventory. The analysis is therefore independent of the overall optimality of the supply schedule. Rather, we focus on how car dealers set prices conditional on inventory levels which are determined by orders placed about 90 days before current sales.

1 Pricing Model

Suppose that a dealer has a lot size of $L > 1$. This determines the maximum number of cars the dealer can hold in inventory at any given time. One customer arrives every period and has a reservation price r drawn from a distribution g_r . The dealer receives a shipment of $S \leq L$ cars every T periods. This supply is fixed in the short run and is thus treated as exogenous for the dealer’s pricing decision.

Holding inventory is costly for the dealer.²⁶ For every car that is in inventory, the dealer incurs inventory cost i per period. If the dealer has no cars on the lot, she cannot sell any cars until the next shipment. We assume that customers drop out of the market or purchase from another dealer if they find no inventory. If the dealer has more than $L - S$ cars on the lot when a supply of S cars arrives, we assume that the dealer must return the cars that do not fit on the lot to the manufacturer and, in so doing, incurs a “return fee” $f \geq 0$ for each returned car.²⁷

We assume that price is determined according to a standard Nash bargaining model. The price paid by the customer who arrives at the dealership at time t (p_t) is a function of the dealership’s opportunity cost (o), the buyer’s reservation price (r), and the bargaining power λ of the seller relative to the buyer. Because exactly one customer arrives each period, we subscript buyers’ reservation prices r and bargaining power λ with t :

$$p_t = (r_t - o_t)\lambda_t + o_t \tag{1}$$

This expression assumes that each party earns its disagreement payoff (what it would earn if negotiations were to fail) plus a share of the incremental gains from trade in time t , with proportion $\lambda_t \in [0, 1]$ going to the seller. When $\lambda_t = 1$, the dealer sells at the reservation price of the buyer. When $\lambda_t = 0$, the dealer has no bargaining power and sells at her opportunity cost.²⁸ To understand

²⁶These costs can be thought of as the “floorplan” cost (the financing cost for cars that are on the lot) and/or as the cost of maintaining the cars in selling condition.

²⁷In practice, dealers can increase their lot size by renting or purchasing more land. The lot size constraint and return fee in the model captures that doing so is costly while constraining the state space of the dynamic model.

²⁸One might think that the bargaining power λ of buyers should vary with the amount of cars in inventory. However, bargaining power and cars in inventory relate to quite distinct concepts. Bargaining power is typically related to the

how inventory and the time until the next shipment affect prices, we must now determine how they affect the opportunity cost of the dealer. Intuitively, the dealer has to trade off selling the car today versus waiting until tomorrow to potentially sell the car. Selling a car tomorrow has the advantage that the dealer might be able to sell the car to a buyer with a higher valuation. The disadvantage is that the dealer will incur the cost of holding the car in inventory for one more period and that the time value of money renders a sale tomorrow worth less than a sale today. To formulate this problem more precisely, we now set up a Bellman equation that describes the dealer's profit as a function of inventory and time relative to when the next shipment arrives. This allows us to specify the opportunity cost of the dealer, o_t , in terms of the dealer's continuation profits for different inventory levels.

Define an inventory cycle c as the set of time periods between two shipments. We number time periods within inventory cycles, i.e., at $t = 1$ a shipment arrives. $t = T$ is the last period of cycle c . Cycle $c + 1$ starts the next period with a new shipment of size S . We can write the dealer's profit in period $1 < t < T$ of cycle c given inventory $n \geq 1$ as:

$$\begin{aligned} \Pi(t, c, n) = & Pr(r_t \geq o_t) (E_\lambda [E_r [\lambda(r_t - o_t) + o_t \mid r_t \geq o_t]] + \delta \Pi(t + 1, c, n - 1)) + \\ & Pr(r_t < o_t) \delta \Pi(t + 1, c, n) - n \cdot i \end{aligned} \quad (2)$$

where $o_t = \delta (\Pi(t + 1, c, n) - \Pi(t + 1, c, n - 1))$ and δ is the discount factor. We assume that the inventory holding cost c accrues at the beginning of each period, i.e., applies to the number of cars n with which the dealer entered the period. To understand the dealer's profit, notice that the dealer will sell a car if the reservation price of the buyer exceeds the dealer's opportunity cost ($r_t \geq o_t$). In this case, the dealer will obtain revenue of $\lambda(r_t - o_t) + o_t$ and enter period $t + 1$ with $n - 1$ cars. If there are no gains from trade ($r_t < o_t$) the dealer will not sell a car and enters period $t + 1$ with n cars. The opportunity cost of the dealer, o_t , is the difference in the dealer's continuation profits from entering the next period with n cars instead of $n - 1$ cars.

At the end of an inventory cycle (period T) the dealer, if she sells a car in period T , enters the next inventory cycle $c + 1$ with $n - 1 + S$ cars on her lot; this is because the dealer receives a shipment of S cars to start the next inventory cycle. If she does not sell a car in the last period of the inventory cycle, she enters the next inventory cycle $c + 1$ with $n + S$ cars. Formally,

$$\begin{aligned} \Pi(T, c, n) = & Pr(r_T \geq o_T) (E_\lambda [E_r [\lambda(r_T - o_T) + o_T \mid r \geq o_T]] + \delta \Pi(1, c + 1, n - 1 + S)) + \\ & Pr(r_T < o_T) \delta \Pi(1, c + 1, n + S) - n \cdot i \end{aligned} \quad (3)$$

where $o_T = \delta (\Pi(1, c + 1, n + S) - \Pi(1, c + 1, n - 1 + S))$.

At the beginning of a new inventory cycle (period 1), the dealer may have to return cars if the shipment S exceeded the available space on the lot when the last inventory cycle ended. In particular, if the dealer entered the new inventory cycle with n cars (including the new shipment S), she needs to return $\max\{0, n - L\}$ cars to the manufacturer at a return fee f per returned car.

degree of patience and/or information asymmetries between the parties. Cars in inventory determine the dealer's opportunity cost o of selling a vehicle, which reflects the dealer's disagreement payoff. See Kennan and Wilson (1993) and Ausubel and Deneckere (1998) for excellent review papers of the relevant game-theoretic bargaining literature.

Inventory cost needs to be paid only for cars that are not returned. Formally,

$$\begin{aligned} \Pi(1, c, n) = & -f \max\{0, n - L\} + \\ & Pr(r_t \geq o_t) (E_\lambda [E_r [\lambda(r_t - o_t) + o_t \mid r_t \geq o_t]] + \delta \Pi(2, c, \min\{n, L\} - 1)) + \\ & Pr(r_t < o_t) \delta \Pi(2, c, \min\{n, L\}) - \min\{n, L\} \cdot i \end{aligned} \quad (4)$$

where $o_T = \delta (\Pi(2, c, \min\{n, L\}) - \Pi(2, c, \min\{n, L\} - 1))$.

To fully characterize dealer profits, if the dealer has no inventory, her continuation profits are the discounted profits of the first period of the new inventory cycle.

$$\Pi(0, t, c) = \delta^{T+1-t} \Pi(S, 1, c + 1) \quad (5)$$

Using (2), (3), (4) and (5), we can derive the dealer's opportunity cost and the expected price for a simple example in which an inventory cycle lasts three periods, the dealer is supplied with exactly one vehicle at the beginning of each cycle, the dealer's lot holds at most three cars, the dealer's per-car and period-inventory cost is 0.001, the discount factor is 0.99, and the dealer's return fee is 0.01 ($S = 1, T = 3, L = 3, i = 0.001, \delta = 0.99, f = 0.01$). Also, we assume that the bargaining power λ and the reservation price of the buyer r are identically but independently distributed uniformly between 0 and 1. We find that, in steady state, the opportunity costs of the seller are as follows:

	Dealer's Opportunity Cost		
	$t = 1$	$t = 2$	$t = 3$
$n = 1$	0.47	0.30	0.18
$n = 2$	0.41	0.26	0.13
$n = 3$	0.34	0.23	-0.0099

To get a sense for how the dealer's opportunity cost changes, first fix an inventory level, for example $n = 1$, and consider the change in opportunity cost as we move closer to the next shipment. A dealer who has one car on the lot in period 1 has two more opportunities to sell that car before she receives a replacement car if she does not sell the car today. Thus, she holds out for a high-valuation buyer by setting the minimum offer she is willing to accept at 0.47. In the next period, the dealer has only one opportunity to sell that vehicle before the next shipment to a buyer who might have a higher reservation price than today's buyer, resulting in a lower opportunity cost for the vehicle. In the third period, the dealer has no other opportunity to sell the car before the next shipment arrives, and the opportunity cost falls still further. It does not fall to zero because the dealer, with two open spaces on the lot, can hold onto the car and sell it in the next inventory cycle. Of course, in doing so the dealer will incur inventory holding cost. One might think that, for $n = 1$, how close the dealer is to the next shipment should not matter; after all, the dealer's lot is large enough to accommodate both the old and the new car. However, holding out too long for a high-valuation buyer increases the probability that during a subsequent inventory cycle the dealer is going to run into a lot-size constraint.

To further gain an understanding how the dealer's opportunity cost changes, now fix the number of periods until the next shipment, for example $t = 3$ (meaning a shipment arrives next period), and consider the change in opportunity cost when the dealer has more cars on the lot. If the dealer has only one car in inventory with a shipment coming the next period, she holds out for a higher

valuation buyer than if she has two cars in inventory. This is for two reasons. First, the cost of holding two cars in inventory is higher than that of holding only one car. In addition, when the dealer has two cars on the lot, she wants to reduce the probability that she will start the next inventory cycle with three cars on the lot—which increases the probability of eventually running into the inventory constraint. Finally, notice that if the dealer will be resupplied next period and has three cars on the lot, her opportunity cost is negative, i.e., the dealer would be willing to accept a negative payment from a customer. This is because for $n = 3$ in the last period before a new shipment ($t = 3$), if the dealer does not sell the car, she will have to pay a return cost of 0.01, which in present value is 0.0099. Hence, the dealer is better off accepting any offer rather than paying the return cost.

Next, we calculate the expected negotiated prices using equation (1):

	Expected negotiated prices		
	$t = 1$	$t = 2$	$t = 3$
$n = 1$	0.57	0.45	0.36
$n = 2$	0.53	0.42	0.33
$n = 3$	0.47	0.39	0.245

The key comparative statics from this example are, first, that holding inventory constant, prices decrease as we move closer to a new shipment. Second, holding the time until a new shipment constant, prices decrease as there are more cars in inventory.²⁹ These comparative static predictions are not unique to this setup; they are shared across a class of models in operations research in which firms face the problem of selling a given stock of items by a deadline, demand is downward sloping and stochastic, and a firm’s objective is to maximize expected revenues (see Yano and Gilbert (2003) and Chen and Simchi-Levi (2012) for detailed reviews of this literature). The existing class of models differ from our setup in two ways: most assume that prices are set by a monopolist instead of being negotiated, and that a given stock has to be sold by a deadline instead of having to sell recurring shipments over an infinite horizon. One of the closest papers to our own in this line of research is a model by Gallego and Ryzin (1994), which characterizes the profit-maximizing prices of a monopolist over a finite horizon as a function of the inventory and the time remaining until the deadline. Their model allows for a salvage value at the end of the (single) inventory cycle and is thus a good representation of pricing within one cycle in our model, except for the fact that their salvage value is linear in the number of units left at the end of the inventory cycle, whereas in our model the value of inventory that carries over into the next period is non-linear in quantity. Kuo, Ahn, and Aydin (2011) present a model that allows customers to initiate a negotiation. However, the remainder of the model does not describe auto retailing well. In their model, retailers sell short, life-cycle durable goods with a fixed selling season and no replenishment. Prices and inventory levels therefore decrease as the season’s end approaches. The point of the paper is that retailers can benefit from allowing some customers to negotiate, especially in cases when the inventory levels are high or the remaining season is short. Another class of models solves versions of the so-called “Knapsack” problem in which an agent has to decide which of stochastically arriving items of different values to include in a “Knapsack” with finite capacity (see, for example,

²⁹These comparative statics hold for all of the many different parameter values for which we have solved this model.

Papastavrou, Rajagopalan, and Kleywegt (1996)). These models also yield the same comparative statics as our example.

Finally, there are four important features of the model to note. First, while the model makes a clear prediction that, holding the time until a new shipment constant, prices decrease as there are more cars in inventory, the model yields no general prediction about the relative *size* of the inventory effect over different days until the next shipment. Similarly, while the model predicts that, holding inventory constant, prices decrease as we move closer to a new shipment, it generates no general predictions about whether this effect is larger for small or large inventories. Hence, the existence and direction of such interactions will be an empirical question.

Second, in steady state, the dealer is in each inventory state with a reasonable probability, except for the full capacity state which happens extremely rarely.³⁰ This will be important for identifying the price effects of inventory in our empirical analysis.

	Steady state probabilities		
	$t = 1$	$t = 2$	$t = 3$
$n = 0$	0	0.517	0.844
$n = 1$	0.968	0.469	0.152
$n = 2$	0.031	0.013	0.004
$n = 3$	0.0005	0.0002	0.00004

Third, the price effect occurs despite the fact that the dealer is correct about the distribution from which the reservation prices of buyers are drawn. In other words, price fluctuations by the dealer are *not* the result of the dealer updating her expectation or “learning” about the underlying level of demand. Price changes occur simply because the dealer balances stochastic demand and fixed short-run supply.

Fourth, the price effect of inventory and time to next shipment are not dependent on whether shipment quantities and/or inventory cycles are chosen optimally by the dealer. The hypothesized effects arise because, regardless of how shipment quantities and/or inventory cycles are chosen, demand is stochastic while supply from the dealer’s perspective is fixed at least 45 days before cars are delivered and typically 90 days before.

Of course there are reasons for a dealer to alter the price at which she sells a car, other than those factors featured in the model. We do not include these other reasons in the formal model for parsimony and simplicity. However, we will consider and include measures of such concepts in the empirical work.

2 Instrumental variables

In this section we use an instrumental variables approach to estimate the effect of inventory on price levels. This technique is more general than our fixed effect specifications in that it will produce consistent estimates in the presence of any form of unobserved demand shock, not just those we described in the previous section.

³⁰Because the firm is resupplied in period 1, it never has 0 cars in period 1.

We need an instrument that is correlated with dealer inventory levels but is uncorrelated with demand shocks that may affect price levels. As in Bresnahan and Ramey (1994), we use *exogenous* plant closures in the US, Mexico, and Canada to construct such an instrument. We exclude plant closures that occurred because demand was weaker than expected. Such closures are intended to prevent inventory build-ups and are correlated with demand shocks that may affect price levels. Our data on plant closures come from *Automotive News*, a trade publication that lists every plant closure in the US, Mexico, and Canada, the duration of the closure, and the reason the plant was shut down. Based on these data we classified the following reasons for plant closures as exogenous:

Reason for plant closure	Closure plant-days
Design problem	76
Earthquake	108
Engine shortage	12
Explosion at parts plant	56
Faulty control arm	8
Faulty cooling system	15
Fire	1
Firestone tire shortage	154
Fix faulty equip	112
Parts shortage	368
Snowstorm	44
Storm or Flood	138
Strike at supplying parts plant	2213
Terrorist attacks	135

This table also contains the number of total plant-days of closure due to the different reasons we classified as exogenous.³¹ Note that most closures in our data are due to strikes at supplying parts plants. Those are strikes that occur in plants that serve the auto manufacturer plant and not strikes at the manufacturer plant. Because the exogenous closures in the table are likely not anticipated by manufacturers or by dealers, it is plausible that they are correlated with dealer’s inventory level but uncorrelated with demand shocks, and therefore they satisfy the relevance and the exclusion restriction.

Because we do not know exactly how many days a plant closure might delay inventory receipt at local dealers, we construct a series of variables that code the number of days that the plant producing the transacted car was closed during two-week periods that fall 5–6, 7–8, 9–10, and 11–12 weeks before the date on which the car was sold. Because plant closures can affect different cars differently, depending on how popular they are, we interact these variables with car dummies to create our instruments. We use a 2SLS specification where we instrument for all four inventory relations variables, namely *Inventory (1–14)*, *Inventory (15+)*, *DaysToResupply*, and *DaysToResupply* × *LowInventory*.

We restrict the sample for this specification to models that were produced in a plant that closed for one of the reasons listed above. Because we only observe plant closures in North America, this restricts the data set to cars produced by the following: Chrysler (149 closure plant-days), Ford (593 closure plant-days), General Motors (2362 closure plant-days), Honda (56 closure plant-days), Isuzu (77 closure plant-days), Mazda (73 closure plant-days), Mercedes-Benz (8 closure

³¹One may suspect that plant closures due to “terrorist attacks” were initiated because manufacturers anticipated weakening demand in the aftermath of the attacks. If this were the case, these plant closures would not be a valid instrument. This, however, is not the case for the plant closures that we have included in our data. Most of these plant closures happened between 9/11/01 and 9/13/01 and seem to have been prompted by a desire not to require workers to come in during the immediate aftermath of the attacks.

plant-days), Mitsubishi (1 closure plant-day), Nissan (20 closure plant-days), Toyota (108 closure plant-days), Volkswagen (26 closure plant-days), and Volvo (4 closure plant-days). This leaves 838,572 observations for the instrumental variables estimation.

We begin by reestimating our standard specification on the smaller dataset. The effect of inventory on price in column 1 of Table A-1 is of slightly smaller magnitude than in the full-sample estimates for below-median inventory levels (-0.036 vs. -0.046) and of slightly higher magnitude for above-median inventory levels (-0.0072 vs. -0.005). The estimates of the IV specification are in column 2 of Table A-1. The IV point estimates on below-median levels of inventory are substantially larger in magnitude than the OLS estimates. In particular, the inventory coefficient for below-median inventory levels is -0.2 (with a p-value < 0.001) vs. -0.036 while the inventory coefficient for above-median inventory levels is slightly smaller than the OLS estimate, at -0.0052 (compared to -0.0072), with a p-value of 0.047.

Our IV estimation is necessarily limited because, although our instrument is clearly exogenous, it is also relatively coarse: the reason is that we are using an instrument (plant closures) that applies to all dealers in our sample equally in order to predict the dealer-specific inventory for a car. Our IV estimates should thus be considered only supporting evidence for the negative effect of inventory on price that we have found persisting across a number of different fixed effect models. We thus continue the paper using a fixed effect specification, which allows us to use the entire dataset.

Table A-1: Basic Results: Instrumental variables[†]

Dep. Var. ln(price)	(1) OLS (on IV sample)	(2) IV (on IV sample)
Inventory (1–14)	-.036** (.0028)	-.2** (.019)
Inventory (15+)	-.0072** (.00049)	-.0052* (.0026)
DaysToResupply	.00043 (.00077)	.012 (.0085)
TradedCar	-.24** (.023)	-.32** (.029)
Tradein	2.8** (.015)	2.8** (.015)
VehicleCost	83** (.27)	83** (.14)
Model Age 5–13 Months	.16** (.03)	.25** (.032)
Model Age > 14 Months	-.038 (.056)	-.023 (.055)
Weekend	.1** (.016)	.12** (.019)
EndOfMonth	-.51** (.017)	-.48** (.018)
EndOfYear	-.19** (.05)	-.2** (.052)
Observations	838,572	838,572
Adj. R-squared	0.951	0.357

[†] Significant at 10%; [*] significant at 5%; ** significant at 1%. Robust standard errors in parentheses.

[†] Unreported are the constant term, car*dealer, monthly fixed effect, and the demographic variables reported in Table 1. The reported R-squared for the IV regression is the within R-squared.

All coefficients are multiplied by 100.