

NBER WORKING PAPER SERIES

SCARCITY RENTS IN CAR RETAILING: EVIDENCE
FROM INVENTORY FLUCTUATIONS AT DEALERSHIPS

Florian Zettelmeyer
Fiona Scott Morton
Jorge Silva-Risso

Working Paper 12177
<http://www.nber.org/papers/w12177>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2006

We are grateful for helpful comments from Severin Borenstein, Dennis Carlton, Jose Silva, Candi Yano, participants at the NBER IO Program Meeting, the Marketing in Israel Conference, seminar participants at UC Berkeley, and at the University of East Anglia. We are especially grateful to Meghan Busse for extensive comments and to Thomas Hubbard for the suggestion to look into inventories. In an early draft this paper was titled “Inventory Fluctuations and Price Discrimination: The Determinants of Price Variation in Car Retailing.” Addresses for correspondence: School of Management, Yale University, PO Box 208200, New Haven CT 06520-8200; Haas School of Business, UC Berkeley, Berkeley CA 94720-1900; Anderson Graduate School of Management, UC Riverside, Riverside, CA 92521. E-mail: florian@haas.berkeley.edu, fiona.scottmorton@yale.edu, jorge.silva-risso@ucr.edu. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2006 by Florian Zettelmeyer, Fiona Scott Morton and Jorge Silva-Risso. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Scarcity Rents in Car Retailing: Evidence from Inventory Fluctuations at Dealerships
Florian Zettelmeyer, Fiona Scott Morton and Jorge Silva-Risso
NBER Working Paper No. 12177
April 2006
JEL No. L0, L1

ABSTRACT

Price variation for identical cars at the same dealership is commonly assumed to arise because dealers with market power are able to price discriminate among their customers. In this paper we show that while price discrimination may be one element of price variation, price variation also arises from inventory fluctuations. Inventory fluctuations create scarcity rents for cars that are in short supply. The price variation due to inventory fluctuations thus functions to efficiently allocate particular cars that are in restricted supply to those customers who value them most highly. Our empirical results show that a dealership moving from a situation of inventory shortage to an average inventory level lowers transaction prices by about 1% ceteris paribus, corresponding to 15% of dealers' average per vehicle profit margin or \$250 on the average car. Shorter resupply times also decrease transaction prices for cars in high demand. For traditional dealerships, inventory explains 49% of the combined inventory and demographic components of the predicted price. For so-called 'no-haggle' dealerships, the percentage explained by inventory increases to 74%.

Florian Zettelmeyer
Haas School of Business
UC, Berkeley
Berkeley, CA 94720
and NBER
florian@haas.berkeley.edu

Fiona Scott Morton
Yale School of Management
135 Prospect Street
New Haven, CT 06520
fiona.scottmorton@yale.edu

Jorge Silva-Risso
Anderson Hall 11
University of California
Riverside, CA 92521
jorge.silva-risso@ucr.edu

1 Introduction

Why do identical cars at the same dealership sell for different prices to different customers? At a superficial level, the reason is that prices are individually negotiated between dealers and customers. The obvious follow-up question is, why are retail automobile prices individually negotiated? A common answer is that negotiation is a way for dealers with market power to price discriminate among their customers. Given the high price of a new car, it would not be surprising if the cost of gaining information about a consumer's willingness to pay is, in comparison, small enough to make the dealer's effort to assess a consumer's valuation and negotiate individual prices more profitable than posting a fixed price.

In this paper we argue that price discrimination is not the only reason why car prices are negotiated. Because car supply is restricted in the short term to the inventory on a dealer's lot, and demand is volatile, the opportunity cost of selling a car of a specific make, model, options, and color is constantly changing with demand for that particular car within the geographic market. Even 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. Thus there are effectively new, dealer-level optimal prices each day - or perhaps more frequently - for each car. Not posting a price, and instead negotiating with the consumer, allows the dealer to 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—is likely to vary across two consumers who purchase the same car on different days, even without differences between them in willingness to pay or bargaining ability. This explanation for price variation differs importantly from the price discrimination explanation because it does not imply that there is market power. Indeed, under the inventory explanation, price differences are the result of scarcity rents, and function to efficiently allocate particular cars that are in restricted supply to those customers who value them most highly.

While these two explanations are very different, they are also not mutually exclusive. For example, there is no reason that a dealer with market power would not vary its price both according to the willingness to pay of individual customers and according to the opportunity cost of the vehicle induced by inventory scarcity (see Borenstein and Rose (1994) for an example of this behavior in the airline industry). The purpose of this paper is to argue that inventory scarcity may be an important but neglected component of price variation, and to estimate the extent to which inventory concerns can explain the variation in prices in retail automobile sales.

We construct a simple dynamic model of a car dealer's pricing problem as a function of inventory. Solving this model for particular parameter values, we find that the price a dealer charges should vary with two factors: the amount of inventory of that specific car in his lot

and the number of days remaining until a shipment of new inventory arrives.

The intuition behind the effect of inventory on prices is as follows. If a dealer’s inventory of particular car is increased, *with no change in the 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, any sale when inventory is very low has a higher opportunity cost because the dealer may not be able to sell to a future high-valuation consumer who could arrive after the last car is sold but before the new inventory arrives. To understand the intuition behind the effect of resupply time on prices, consider a dealer who is approaching the date when a new shipment of a particular car will arrive. As the date nears, the opportunity cost of selling the remaining cars on his lot falls *conditional on the inventory level*, because soon the dealer will be restocked. Thus, as days to resupply falls, the dealer will be more willing to discount the car to a consumer with a low valuation. As we will show later, the price effect of inventory occurs despite the fact that the dealer is correct about the distribution from which the reservation prices of buyers are drawn; the dealer is not updating his expectation or “learning” about the underlying level of demand.

Our model is related to a known class of models in the operations research literature which relate prices to inventory in a monopoly pricing environment (see Yano and Gilbert (2003) for a detailed review of this literature). Our model shows that it is possible for these results to carry over to a negotiated price environment.

The empirical section of the paper provides evidence for the relationship of prices to both inventory levels and resupply times. A dealership moving from a situation of inventory shortage to an average inventory level lowers transaction prices by about 1% *ceteris paribus*, corresponding to 15% of dealers’ average per vehicle profit margin or \$250 on the average car. Additionally, shorter resupply times decrease transaction prices for cars in high demand. We consider the potential endogeneity of prices and inventory levels due to, for example, a temporary demand shock that raises the price of a model and lowers inventory levels. We use a series of fixed effects specifications as well as instrumental variables to control for this potential problem. Our results remain robust to these approaches, as well as to alternative definitions of inventory. We present some extensions of our results showing that local inventory also affects transaction prices and that the effect of inventory is stronger for car models in high demand. Finally, we use our estimates to calculate that the share of the price variation attributable to either inventory or demographics that is due to inventory is 49% in our sample. For so-called “no-haggle” dealerships, inventory explains 74% of the combined inventory and demographic components of the price variation.

We are not aware of any empirical work in economics that discusses inventory fluctuations as

a source of scarcity rents. However, there is some research that analyzes the interplay of prices and inventory, albeit with a substantially different focus than our paper. In particular, there are recent papers that solve a dynamic model of optimal inventory investment and estimate it with data on prices and inventories. Hall and Rust (2000) build such a model to 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 (2005) model the optimal pricing and production decisions of auto manufacturers which sell overlapping vintages of the same product simultaneously. They estimate their model using aggregate data on transaction prices, quantities, and inventories, and find that prices of cars fall by 9% over the course of a model year. A key conclusion of the paper is that about 1/2 of this price decline is driven by a ‘build-to-stock’ (as opposed to a ‘build-to-order’) policy practiced by the manufacturers. Copeland and Hall (2005) examine how the Big Three auto makers accommodate shocks to demand. They estimate a dynamic profit maximizing model of the firm that takes inventories into account and show that when a manufacturer is exposed to a demand shock, sales adjust immediately, prices adjust gradually, and production adjusts only after a delay.

The paper proceeds as follows. In section 2 we develop a simple model to illustrate the relation of prices and inventory levels and derive empirical predictions. In section 3 we describe our data and discuss the measurement of inventory in that context. In section 4 we discuss estimation issues. In section 5 we estimate the price-inventory relationship. In section 6 we analyze the robustness of the empirical results. In section 7 we consider a number of extensions to the basic results. In section 8 we determine the share of the price variation attributable to either inventory or demographics that is due to inventory. In section 9 we conclude.

2 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 an example. We use the insights from the model to derive empirical predictions.

2.1 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 consumer 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. We explored whether this assumption is an accurate reflection of the supply relationship between dealers and manufacturers. In interviews with car dealers and manufacturers we found that, while dealers order frequently from manufacturers, the lag between the order and when the dealer actually receives the car is at least 45 days and typically 90 days. Within that time period, dealers cannot obtain additional cars from the manufacturer for delivery at that shipping date.¹ Also, they cannot reduce their order, or alter its composition.²

If the dealer has no cars on the lot, he cannot sell any cars until the next shipment. We assume that consumers 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 has to return the cars that do not fit on the lot to the manufacturer and in doing so incurs a “return fee” $f \geq 0$ for each returned car.³

We assume that price is determined according to a standard Roth-Nash bargaining model. The price paid by the consumer 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. Since exactly one consumer arrives each period, we subscript consumers’ 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 his opportunity cost.⁴

¹However, they can exchange vehicles with other dealers. We do not consider this possibility in the model, but in the empirical analysis we control for inter-dealer trades. Please see section 3.2 for a discussion of dealer trades.

²Because of our focus on the dealer’s short run pricing problem we not address in this model 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 his lot and then compensating consumers for the benefit or cost of that choice with the price of the car. The model presented in this section will link prices to inventory levels, irrespective of how inventory levels were chosen by the dealer, and will thus apply to the situations discussed in Carlton (1978) and Dana (2001). 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.

³The lot size constraint together with this “return fee” has the same effect in this model as an inventory holding cost (see below).

⁴One might argue that the bargaining power of consumers should vary with the amount of cars in inventory. However, this assumes that consumers are aware of inventory levels. Since dealers typically store a large fraction

In order to understand 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 and selling the car to a buyer who might have a higher valuation. 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:

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

where $o_t = \Pi(n, t + 1, c) - \Pi(n - 1, t + 1, c)$. 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 he sells a car in period T , enters the next inventory cycle $c + 1$ with $n - 1 + S$ cars on his lot; this is because the dealer receives a shipment of S cars to start the next inventory cycle. If the dealer does not sell a car in the last period of the inventory cycle he enters the next inventory cycle $c + 1$ with $n + S$ cars. Formally,

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

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

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 at the end of the last inventory cycle. 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

of cars in lots that are not visible from the front of the dealership, or in separate (cheaper) back lots, consumers are not normally not able to assess how many cars a dealer has in stock at any particular time. Although a few dealers have begun posting inventory in recent years, this was very rare at the time of our sample.

returned car. Formally,

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

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

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

$$\Pi(0, t, c) = \Pi(S, 1, c + 1) \quad (5)$$

Using (2), (3), (4) and (5), we can derive the opportunity cost of the buyer and the expected price for a simple example in which an inventory cycle lasts 3 periods, the dealer is supplied with exactly one vehicle at the beginning of each cycle, the dealer's lot holds at most 3 cars, and the dealer's return fee is 0.05 ($S = 1, T = 3, L = 3, f = 0.05$). 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 cost of the seller are as follows:

	Dealer's Opportunity Cost		
	$t = 1$	$t = 2$	$t = 3$
$n = 1$	0.56	0.50	0.41
$n = 2$	0.36	0.30	0.24
$n = 3$	0.17	0.08	-0.05

To get an intuition 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 he receives a replacement car if he does not sell the car today and thus holds out for a high valuation buyer by setting the minimum offer he is willing to accept at 0.56. In the next period, the dealer has only one opportunity to sell that vehicle before the next shipment to a buyer who may 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 since the dealer, with two open spaces on the lot, can hold on to the car and sell it in the next inventory cycle. One might think that for $n = 1$ it should not matter how close the dealer is to the next shipment since even if the next shipment arrives, 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 understand the intuition for 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 as the dealer has more cars on the lot. If the dealer has only one car in inventory with a shipment coming the next period, he holds out for a higher valuation buyer than if he has two cars in inventory. This is because in the latter case he wants to reduce the probability that he will start the next inventory cycle with three cars on the lot—increasing 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, his opportunity cost is negative, i.e. the dealer would be willing to accept a negative payment from a consumer. This is because for $n = 3$ in the last period before a new shipment ($t = 3$), if the dealer does not sell the car he will have to pay a return cost of 0.05. Hence, the dealer is better off accepting a small negative offer rather than paying the return cost.

Of course, the dealer’s opportunity costs are not the negotiated prices unless $\lambda = 0$. The expected negotiated prices can be derived by taking the expectation over r and λ in equation (1):

	Expected negotiated prices		
	$t = 1$	$t = 2$	$t = 3$
$n = 1$	0.67	0.62	0.56
$n = 2$	0.52	0.47	0.43
$n = 3$	0.38	0.31	0.23

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.⁵ 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) for a detailed review of this literature). The existing class of models differ from our setup in two ways: they 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

⁵These comparative statics hold for all of the many different parameter values for which we have solved this model.

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 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. Another class of models solve 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, 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 full capacity which happens extremely rarely.⁶ This will be important for identifying the price effects of inventory in our empirical analysis.

	Steady state probabilities		
	$t = 1$	$t = 2$	$t = 3$
$n = 0$	0.00	0.37	0.66
$n = 1$	0.85	0.57	0.32
$n = 2$	0.15	0.06	0.02
$n = 3$	0.004	0.0007	0.00006

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 his 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.⁷ The

⁶Since the firm is resupplied in period 1, it never has 0 cars in period 1.

⁷With regards to the optimality of the dealer’s choices, we have anecdotal evidence that manufacturers play a role in what dealers can order, and that supply is not generally chosen in the best interest of the dealer. For

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.

2.2 Empirical predictions

We have several goals for the paper. Our first and primary goal is to see whether our two comparative statics hold:

- Controlling for the time until a new shipment arrives, prices decrease as there are more cars in inventory.
- Controlling for inventory, prices decrease as a dealer moves closer to a new shipment.

A second goal in our paper is to explore some other factors that affect the inventory-price relationship. In particular, the inventory-price relationship may be different for cars which are in overall short supply. For these car models it is close to impossible for dealers to make use of dealer trades to obtain additional cars of a particular type on short notice if their demand realization for that car is higher than expected.⁸ Thus, we expect to observe that inventory effects are larger for these cars than for others where dealer trades may be costly but possible. We also examine whether the price-inventory relationship differs at the end of the month. Manufacturers and dealers impose non-linear sales targets on sales personnel which increase their incentive to sell additional cars at the end of the month. Due to these incentives we expect to observe that at the end of the month dealers are less likely to hold out for a high price when inventory levels are low than they would at other times. We also expect to find a stronger price effect of inventory when a dealer lot is close to maximum capacity. This is because the inventory cost associated with the last few cars which the dealer can store on his lot may be particularly high. For example, the dealer might have to use customer parking, thereby decreasing revenue, or reduce spacing between parked cars, increasing the danger of damaging cars on the lot. Hence, we expect that inventory that is close to lot capacity will decrease the dealer’s opportunity cost of selling a car and thus decrease transaction prices.

The third goal in our paper is to test empirical predictions about how much of the variation in car prices can be explained with price discrimination and how much can be explained with inventory fluctuations. Since we measure only a subset of what dealers observe when price discriminating among consumers, we cannot estimate precisely what fraction of the price

example, manufacturers often force dealers to take delivery of low demand cars as a condition for obtaining some high demand cars.

⁸We discuss dealer trades in detail in section 3.2.

variation can be attributed to these two factors. We can, however, predict how this fraction should change as a function of the type of dealership. Since many consumers have a high disutility of bargaining, we have seen over the last decade the emergence of dealerships which promise consumers a “no-haggle” price. Most prominently, this is true for the dealership chain AutoNation. While the “no-haggle” policy is popularly believed to mean “fixed-price” this is not true; these dealerships set car prices daily based on their inventory and demand conditions. The no-haggle policy simply means that salespeople are discouraged from varying price across buyers who arrive on a particular day. Consequently, we expect that inventory fluctuations should explain a larger percentage of price variation for AutoNation dealerships than other dealerships.

3 Data

Our data come from a major supplier of marketing research information (henceforth MRI). MRI collects transaction data from a sample of dealers in the major metropolitan areas in the US. We have data containing every new car transaction at California dealerships in the MRI sample from July 1, 1998 to May 31, 2003. These data include customer information, the make, model and trim level of the car, financing information, trade-in information, dealer-added extras, and the profitability of the car and the customer to the dealership.

3.1 Inventory measurement

We measure inventory in our data on the level of the interaction of make, model, model year, body type, transmission, doors, and trim level. 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 1999 or 2000 model, whether it is an EX or LX trim level, whether it is manual or automatic, etc. Tracking inventory on the level of this definition is important because consumers 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 precisely, however, we are making an assumption that consumers essentially do not substitute between versions of a car very easily. We will test this assumption later in the paper.

Since our data are derived from a record of transactions, we do not have a direct measure of inventory. However, we know for every car that was sold how long the car was on the lot. 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 backwards, 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 have been sold during our sample period of four years, thereby generating an observation which allows 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 came on the lot if they get sold during our sample period. Many cars which come onto the lot at the end of our sample period are sold after our observations end. 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 (99.4%) cars fall within this time frame. Hence, our final dataset comprises car purchases for almost four years from July 1, 1998 to May 31st, 2002. Figures 1, 2, and 3 show the inventory levels over time for a Honda, Chevrolet, and Mercedes dealer, respectively. For each of these dealers we have 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 (0-80 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 see this notice that even small dealers need a certain number of cars on the lot to be able to offer variety to consumers. This implies that a large dealer does not necessarily need more cars on the lot compared to 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 0 for small dealerships (see the top panel of Figure 4 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 11 and one for 11 and above. Specifically, we restrict the sample to dealership-car combinations for which the dealership sells at least 3 cars per month according to our definition of a car (see the bottom panel of Figure 4 for a histogram of daily inventories for such dealership-car combinations) and then simply count the cars in inventory. In choosing this third approach we assume that an additional unit of inventory has the same effect for dealers of different size (in section 5.2 we will show empirically that this assumption is a good approximation of reality.) This approach leaves 351,916 observations and will be the sample used throughout the paper. Summary statistics for the dataset are in Table 1.

3.2 Resupply measurement

Since our predictions on inventory are conditioned on the number of days until the next shipment of a car arrives, we need a measure of “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 consumer wants a specific car and the dealer offers to obtain this car for the consumer at another dealership in the region. According to industry participants we interviewed, such “trades” are indeed always an exchange. If the competing dealer agrees on the trade, an employee of the requesting dealership drives an agreed-upon exchange vehicle to the other dealership and brings the requested vehicle back. If the cars are of different value, dealers settle the difference at invoice prices.⁹

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 during 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 which have 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

⁹In 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.

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. Since 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 4 days of arriving on the lot and has an odometer reading that exceeds the 95th quantile as derived above. Since there is a received trade for every requested trade, 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 on that day. Approximately 8% of vehicles are classified as *TradeRequested* and another 8% are classified as *TradeReceived* in the original sample. This matches well with industry estimates that somewhat less than 20% of sold cars are dealer trades.

We can now define *DaysToResupply* as the number of days until a vehicle of the same inventory “car” definition arrives, excluding vehicles that were classified as *TradeRequested* or *TradeReceived*.¹⁰ 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 3 cars per month according to our definition of a car) can be seen in Figure 5. We also graph the distribution of the number of vehicles of the same inventory “car” definition that arrive in a single shipment (see Figure 6).

We will use *TradeRequested* as an indicator variable. Since a dealer bears additional transaction and transportation costs for requested trades, we expect him to pass those on to the consumer.

We have excluded from the data all transaction that fall after 45 days before the introduction of the next model year. We omit these transactions from the dataset as 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.

3.3 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

¹⁰A small percentage of observations end up with very high *DaysToResupply* using this procedure. We drop about 4000 observations where the *DaysToResupply* is greater than six months.

the time of sale that contribute to the resale value of the car.¹¹ The *Price* variable we use as the dependent variable is this price, minus the *ManufacturerRebate*, if any, given directly to the consumer, and minus what is known as the *TradeInOverAllowance*. *TradeInOverAllowance* is the difference between the trade-in price paid by the dealer to the consumer 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 consumers 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.

3.4 Controls

We include a car fixed effect for each interaction of make, model, body type, transmission, displacement, doors, cylinders, and trim level.¹² We drop any observations of “cars” according to this definition with fewer than 200 sales in California during the sample period. Cars with this few sales have hardly any variation in inventory levels. Hence, they are unhelpful in identifying inventory effects but use up degrees of freedom. While 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 by options, we include as an explanatory variable the percent 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 percent 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 to other examples of the same “car” (as specified above).

To control for time variation in prices, we define a dummy *EndOfMonth* that equals 1 if the car was sold within the last 5 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 and also for the dealer’s opportunity

¹¹Dealer-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.

¹²This 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.

cost of not selling the car. 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 later and assign a dummy variable to each category.

We also control for the income, education, occupation, and race of buyers by using census data that MRI matches with the buyer’s address from the transaction record. The data is 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 1100 people in them.

Finally, we control for the region 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.

4 Estimation issues

Conceptually, the bargaining power λ_t of the seller relative to buyer t from equation (1) is quite distinct from the outside options of the seller and the buyer. While the reservation price of the buyer and the opportunity cost of the seller determine the size of the gains from trade, λ_t specifies how the gains from trade are split between the parties. Empirically, we can find measures that are related only to the opportunity cost of the dealer, for example, current inventory and days to resupply. We can also find some measures that are uniquely related to a consumer’s reservation price, for example, the degree of competition between dealers, or the availability of substitutes for the vehicle in question (other brands, models, options, etc.). Nonetheless, our data will not allow us to separately identify a consumer’s relative bargaining power from her reservation price. This is because we have no direct measures of bargaining power, such as patience or the inherent utility or disutility of bargaining for a consumer. Also, the bargaining ability of a buyer may be correlated with measures which also determine a consumer’s reservation price, for example income, educational status, and whether or not a consumer has a car to trade-in. As a result, we will estimate an empirical model in which we will be able to separate price variation due to inventory fluctuations from price variation due to price discrimination. Whether the latter part of price variations is due to heterogeneity in reservation prices or in bargaining abilities is not a question we will be able to answer.

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.

5 The price-inventory relationship

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:

$$\ln(\text{Price}_i) = X_i\alpha + D_i\beta + I_i\gamma + \epsilon_i \tag{6}$$

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 dealer trade indicator.

5.1 Basic specification

To estimate the effect of inventory on prices, we estimate a specification that is informed by the dynamic programming model of section 2.1. The model indicates that prices should increase in days to resupply, controlling for inventory and 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 2-part spline in our specification. In particular, we estimate a different inventory coefficient for below- and above-median inventory levels (the median is 10). Similarly, we also estimate days to resupply as a 2-part spline (split at the median which is 4 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 coefficient. For below median inventory levels (10 and fewer cars), one additional car in inventory is associated with a price that is lower by 0.085% (see variable *Inventory (1-10)*). For above median inventory levels (11 and more cars), one additional car in inventory is associated with a price that is lower by 0.018% (see variable *Inventory (11+)*). An increase in inventory from 1 car to 16 cars (a one standard deviation increase) is associated with a 0.92% reduction in average price. This corresponds to 14% 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 11 to 26 cars is associated with a 0.27% lower average price. This corresponds to 4% of the average dealer gross margin.

Both “days to resupply” variables have an insignificant effect on transaction prices in this specification. We will explore possible reasons for this in the next section. Consumers do not pay a different price when for a vehicle which was requested from another dealership (*TradedCar*). Other coefficient estimates have the expected signs. For example, cars that are sold at the end of the month (*EndOfMonth*) when sales people are trying to meet sales quotas sell for on average 0.31% lower prices. In another example, the prices of cars that are sold more than 4 months after the model was introduced sell for on average 0.98% less than cars sold in the first 4 months of the model cycle. Demographic variables also have the expected sign. For example, women pay slightly more for a car, as do consumers who live in neighborhoods with a higher percentage of residents who have less than high school education. Higher income is associated

with lower prices, except for the highest income consumers.¹³

5.2 Is it reasonable not to scale inventory by dealer size?

Recall from our discussion in Section 3.1 that our inventory variable is the raw number of cars in inventory at each dealer. This implicitly assumes that one additional unit of inventory has the same effect on price irrespective of dealer size. We now investigate whether this assumption is reasonable. To do so we split the sample into 4 subsamples, according to quartiles of the number of vehicles sold by each dealer of a particular car type during the model year: We first calculate for each dealership-car pair the number of vehicles that were sold during the model year. Second, we split the dealership-car pairs into quartiles, according to the sales volume we calculated. Third, we assign each observation in the dataset to the corresponding dealership-car quartile. Finally, we use the quartiles to split the dataset into four subsamples such that the first subsample contains the lowest dealership-car sales volume whereas the fourth subsample contains the highest dealership-car sales volume. For each sample we estimate our basic specification, however, instead of using a spline for inventory we estimate the price effect of inventory non-parametrically by using 10 dummies, one for each decile of inventory as determined from the full dataset. Importantly, this means that the deciles are common across all dealerships and are invariant to dealer size. Thus, each inventory dummy captures the same absolute amount of inventory across all 4 regressions on the different subsamples. If the effect of one additional unit of inventory is the same for small and large dealership-car combinations, we should find that the coefficient of each dummy is approximately the same across all four size quartiles. In addition, if the full sample spline specification in column 1 of Table 2 is correctly specified, we expect that its estimates predict the same absolute price-inventory relationship as is predicted by the non-parametric estimates from the regressions on the four subsamples.

We compare the coefficient estimates for the dummies by graphing the effect of inventory on price across the 10 dummies (see Figure 7). In addition, the graph contains the effect of inventory on price predicted by the full sample spline specification. As can be seen from this graph, the effect of inventory on price for inventories below 10 is extremely similar regardless of dealership-car sales volume. In addition, the full sample spline estimate tracks the non-parametric estimates very well. For inventories above 10 the effect of inventory on price for inventories is also similar for all but the smallest dealership-car sales volume quartile. The reason for the anomalous estimates for the smallest dealership-car sales volume quartile is simple: since there are few cases of very high inventories for the smallest dealership-car sales

¹³For a thorough analysis of the effects of demographics on car prices please see Scott Morton, Zettelmeyer, and Silva-Risso (2003)

volume quartile, there are few high-inventory observations in size quartile 1 and the estimates are thus imprecise.¹⁴ To elaborate, in Figure 8 we repeat Figure 7 for each size quartile separately and with confidence intervals in addition to point estimates. The figure shows that the effect of inventory on price is estimated quite precisely across all four size quartiles for inventories below 10. For inventories above 10, however, the estimates for size quartile 1 become very imprecise. Similarly, for inventories above 20, the estimates for size quartile 2 also become imprecise.

We conclude that using the raw number of cars as our inventory variable is reasonable. This is because the results from the size quartile subsample suggests that one additional unit of inventory has approximately the same effect on price regardless of dealership-car sales volume. In addition, this analysis suggest that our full sample spline specification in column 1 of Table 2 captures the non-linear nature of the price-inventory relationship quite well.

6 Robustness Checks

We now explore the robustness of the estimated price-inventory relationship. We would like to make sure that our results are not due to a potential endogeneity of prices and inventory levels due to demand shocks. In a first approach we use a sequence of fixed effects to address the potential endogeneity of price and inventory (see section 6.1). Second, we use an instrumental variables approach to estimate the effect of inventory on price levels (see section 6.2).

In addition to a potential endogeneity of prices and inventory levels we would also like to determine whether the estimated price-inventory relationship is robust to the level on which inventory is measured. First, we want to make sure that an individual dealer’s inventory for a car does not simply proxy for the local inventory of that car (see section 6.3). Second, we want to make sure that our estimates are not biased by the granular definition of a car we use for constructing our inventory measure (see section 6.4).

6.1 Are the findings due to common demand shocks?

In the next two specifications we repeat the basic specification in 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 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

¹⁴For example, there are only 37 observations to identify the highest inventory deciles in size quartile 1. See Figure 9 for the distribution of observations across size quartiles and inventory dummies.

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.¹⁵ 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 (10 and fewer cars), one additional car in inventory is associated with a price that is lower by 0.092% (see variable *Inventory (1-10)*). For above median inventory levels (11 and more cars), one additional car in inventory is associated with a price that is lower by 0.016% (see variable *Inventory (11+)*). Both coefficients remain precisely estimated despite a substantial decrease in degrees of freedom: while the specification in column 1 contains 1271 car fixed effects and 741 dealer fixed effects, the specification in column 2 contains 8881 car*dealer fixed effects. As before, the days to resupply variables are not statistically significant from zero.

Our second alternative specification accounts for the possibility that demand shocks are short lived *and* local. So far, our monthly fixed effects absorb the price effect of short term demand shocks but only if these affect all vehicle segments in all markets equally. This may not be a good assumption: for example, suppose that a particularly snowy January in the California Sierras increases demand for SUVs for the rest of the winter in the Sacramento area (but not in Southern California), thus simultaneously causing high prices and low inventories for the SUV segment in Sacramento dealerships for that quarter. To rule out this alternative explanation, in column 3 of Table 2 we repeat the specification of column 2 of Table 2 expanding the month fixed effects to month-local area-vehicle segment fixed effects. The local areas are defined as DMAs, metropolitan areas that correspond to TV markets (e.g. Los Angeles, Santa Barbara-San Marino-San Luis Obispo, San Diego, etc.).¹⁶ 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 8881 car*dealer fixed effects and 2871 month*segment*DMA fixed effects (see column 3 of Table 2). We find that for below

¹⁵Of course, a competent dealer in this situation would try to adjust his inventory in the long run and so this story really only applies if this proves difficult or the shock is transitory (see below).

¹⁶Our data contains 12 such local markets: Bakersfield, Chico-Redding, Eureka, Fresno-Visalia, Los Angeles, Monterey-Salinas, Palm Springs, Sacramento-Stockton-Modesto, San Diego, San Francisco-Oakland-San Jose, Santa Barbara-San Marino-San Luis Obispo, and Yuma-El Centro.

median inventory levels (10 and fewer cars), one additional car in inventory is associated with a price that is lower by 0.064% (see variable *Inventory (1-10)*). For above median inventory levels (11 and more cars), one additional car in inventory is associated with a price that is lower by 0.013% (see variable *Inventory (11+)*). Both variables remain precisely estimated. As before, the days to resupply variables are not statistically significant from zero. 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.¹⁷

Before we further explore a potential endogeneity of the price-inventory relationship with an instrumental variables approach, we first would like to better understand why the estimates in Table 2 have thus far not shown the positive relationship between price and days to resupply that the dynamic programming model in Section 2.1 hypothesized. We now explore two reasons for why we might not have found evidence of this relationship.

The first reason could be that days to resupply as measured is too specific to the model setup. The model assumed that resupply occurred at fixed intervals and in fixed quantities. A brief glance at Figures 1, 2, and 3 show that dealers receive cars at irregular intervals and in varying quantities. This suggests that the mechanism by which “days to resupply” and price are related in our data is better described by a variable which measures the number of cars that are due to arrive within set intervals. Hence, we repeat the basic car*franchise fixed effect specification with four variables in addition to *DaysToResupply*. They measure the number of cars of the same type as the transacted vehicle that will arrive on the dealer’s lot one, two, three, and four weeks, respectively, after the transaction. As can be seen in column 1 of Table 3, none of the coefficients of the new variables are significantly different from zero.

The second reason for why we might not have found evidence of a price-days to resupply relationship is that its magnitude depends on the interaction between inventory levels and days to resupply. In particular, it might be that more days to resupply is associated with higher prices in a significant and measurable way only if inventory levels are very low. If inventory levels are low, it is more likely that a dealer that has to wait longer for new supply will run out of inventory before the new shipment arrives. Since this raises the dealer’s opportunity cost of selling a car, we expect to see that for low inventory levels more days to resupply is associated

¹⁷These results are robust to a variety of other fixed effects specifications. For example, we have 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 (1271 dummies) and week*subsegment*DMA fixed effects (19176 dummies). We find that for below median inventory levels (10 and fewer cars), one additional car in inventory is associated with a price that is lower by 0.059% (see variable *Inventory (1-10)*). For above median inventory levels (11 and more cars), one additional car in inventory is associated with a price that is lower by 0.011% (see variable *Inventory (11+)*). Both variables remain precisely estimated.

with higher prices. To investigate this explanation, we interact the variable *DaysToResupply* with an indicator *LowInventory* that equals 1 if inventory levels are in the lowest 10th percentile of average inventories (3 or fewer cars according to our definition are on the lot). The results are reported in column 2 of Table 3 and provide support for the hypothesis. As before, the estimated coefficient on *DaysToResupply* alone is not statistically significantly different from zero, however, when interacted with *LowInventory*, more days to resupply is associated with higher prices (p-value 0.07). Based on this result we use the specification in column 2 of Table 3 as the basis for further analysis.

6.2 Instrumental variables

We now proceed to using 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.

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 why the plant was shut down. Based on this data we classified the following reasons for plant closures as exogenous:

Reason for plant closure	Closure plant-days
Design problem	8
Engine shortage	8
Explosion at parts plant	16
Faulty control arm	12
Faulty cooling system	16
Firestone tire shortage	60
Fix faulty equip	18
Flood	80
Parts shortage	112
Snowstorm	2
Terrorist attacks	127
Strike at plant	45
Strike at supplying parts plant	740

This table also contains the number of total plant-days of closure due to the different reasons we classified as exogenous.¹⁸ Notice that most closures in our data are due to strikes at supplying

¹⁸One may suspect that plant closures due to “terrorist attacks” were initiated because manufacturers antici-

parts plants. Because we do not know the exact lag with which a plant closure affects local dealer inventories we construct a series of variables that code the number of days that the plant that produced the transacted car was closed during two-week periods 5-6, 7-8, 9-10, and 11-12 weeks prior to 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 also interact the plant closure variables with whether the car was sold in Northern or Southern California; this is because manufacturers may decide to change deliveries to regions differentially when they change their delivery schedule due to unanticipated plant closures. We use a 2SLS specification where we instrument for all four inventory relations variables, namely *Inventory (1-10)*, *Inventory(11+)*, *DaysToResupply*, and *DaysToResupply*LowInventory*.

We restrict the sample for this specification to models that were produced in one of the plants 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 Chrysler (124 closure plant-days), Ford (297 closure plant-days), General Motors (888 closure plant-days), and Volkswagen (17 closure plant-days). This leaves 75,213 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 4 is of slightly larger magnitude than in the full-sample estimates for below median inventory levels (-0.11 vs. -0.088) and of slightly smaller magnitude for above median inventory levels (-0.013 vs. -0.016). The estimates of the IV specification are in column 2 of Table 4. The IV point estimates on inventory are similar to the OLS estimates. In particular, the inventory coefficient for below median inventory levels is -0.17 vs. -0.11 while the inventory coefficient for above median inventory levels is -0.012 vs. -0.013. The inventory coefficient for below median inventory levels is highly significant (p-value 0.007), however, the standard error on the inventory coefficient for above median inventory levels is too high to conclude that it is different from zero. Finally, in the IV specification the coefficient on *DaysToResupply*LowInventory* is positive and significant, although, at 0.052 it is larger than in all prior OLS specifications, where it ranged from 0.0032 to 0.0045.

Our IV estimation is necessarily limited because while 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 to predict the dealer-specific inventory for a car. Our IV estimates should thus be considered only supporting evidence for the negative effect

pated 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 and 9/13 and seem to have been prompted by a desire not to require workers to come in during the immediate aftermath of the attacks.

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 allow us to use the entire dataset.

6.3 Does dealer inventory proxy for local inventory?

So far we have interpreted the estimated inventory coefficients as evidence that, holding days to resupply constant, the low inventory of an individual dealer on a specific car induces that dealer to hold out for a higher price than when his inventory is plentiful. In the previous sections we have shown that this finding cannot be attributed to a variety of demand shocks that would simultaneously lead to low inventories and high prices. One explanation we have so far not ruled out is that car-specific dealer inventories in a local area are highly correlated and that therefore, an individual dealer's inventory for a car simply proxies for the local inventory of that car. If that were the case, then our estimate of the price-inventory relationship would be evidence that overall dealer pricing is related to local market conditions, not that the low inventory of an individual dealer on a specific car induces that dealer to hold out for a higher price.

To rule out this explanation we repeat our main specification from column 2 of Table 3 while controlling for the local inventory of the transacted car. We calculate the local inventory by summing inventory within a "car," across all dealers in a DMA. This measure excludes the focal dealer's own inventory. The results are in column 3 of Table 3. We find that the coefficient estimates of the inventory and days to resupply variables are slightly smaller in magnitude than in column 2. While controlling for regional inventory does not change our interpretation that dealers price in reaction to *their own* inventory levels, the coefficient on *LocalInventory* shows that transaction prices are also associated with overall local inventory levels. Consumers who purchase a car when regional inventory is low pay higher prices. The magnitude of the estimated coefficient indicates that a one standard deviation decrease in regional inventory (86 cars) is associated with a price increase of 0.42%.

This analysis suggests that our inventory measure is indeed picking up the effect of scarcity at the dealership level, not at the region.

6.4 Is our inventory measure too narrowly defined?

We have so far measured inventory based on a very granular definition of a car. This may lead us to underestimate the effect of inventory on prices if consumers consider "cars" for which we count inventory separately to be close substitutes: dealers may not have an incentive to hold out for a higher price even if a particular car is in low inventory if consumers are willing to

purchase similar cars for which there is more inventory.

We analyze whether our granular inventory definition affects our results in two ways. First, we see whether the estimated price-inventory relationship changes when we control for the inventory of car variations that are potential substitutes for the car under consideration. The inventory measure we use throughout the paper is at the level of the interaction of make, model, model year, body type, transmission, doors, and trim level. We now define *SubstituteInventory* as the inventory of all cars that share the make, model, model year, body type, and doors of the transacted cars. In other words, this variable capture how many cars the dealer has in inventory that are of the same model, year, and body type (sedan, coupe, convertible, etc.) but differ in transmission and trim level. We now estimate our basic specification with the addition of controlling for *SubstituteInventory*. The results in column 4 of Table 3 show that the estimated price-inventory relationship is essentially unchanged if we hold constant the number of “substitute” cars that are in the dealer’s inventory. We find that the inventory of these substitute cars is negatively related to the transaction price of the car, however, the effect is very small. For each additional substitute car in inventory, the transaction price of the car under consideration decreases by 0.0055%.

Our second approach to analyzing whether our granular inventory definition affects our results is to change the level at which we define inventory. Instead of defining *SubstituteInventory*, we directly redefine our inventory measures at the level of the interaction of make, model, model year, body type, and doors. The results in column 5 of Table 3 show that the estimated price-inventory relationship is slightly smaller under the redefined inventory measure. One additional car in inventory at below median inventory levels decreases price by 0.07% in contrast to 0.09% with the more granular inventory definition. Similarly, one additional car in inventory at above median inventory levels decreases price by 0.0083% in contrast to 0.016% with the more granular inventory definition.

In summary, our basic results seem to be robust to a change in the level at which we measure inventory and to whether we control for the number of “substitute” cars that are in the dealer’s inventory.

7 Extensions

We have predicted in section 2.2 that the inventory-price relationship will be stronger in some cases than in others. In this section we present a series of specifications to test these ideas.

7.1 Does inventory matter differently for cars with strong demand?

One could argue that our stylized model best describes cars that are in strong demand relative to supply. For these car models it is close to impossible for dealers to make use of dealer trades to obtain additional cars on short notice if their demand realization is higher than expected. Thus, we expect to observe that inventory effects are larger for these cars than for others where dealer trades may be costly but possible and for which dealers know that a situation of low inventory will very likely be short-lived.

Consequently, we examine whether the inventory effect is different for cars which face robust demand relative to supply. We expect to observe that inventory effects are larger for these cars than for others. We define such “hot” cars using the average time such cars stay on a dealer’s lot across all dealerships in our sample. We average this time interval, known as “days to turn,” across the whole sample for a car in a calendar year. If *DaysToTurn* of a car in a calendar year is in the lowest quartile, we define the car as “hot.” Notice that this definition of “hot” encompasses a full quarter of all the cars in our sample. These are cars that are generally “in demand,” not necessarily cars that garner media attention for having long waiting lists. This latter group is a much smaller fraction of new cars sold.

We replicate our basic specification (column 2 in Table 3) using interactions of our inventory variables and *DaysToResupply* with the “hot” indicator. The results in column 1 of Table 5 are partially consistent with our hypothesis. While one additional car in inventory at *above* median inventory levels has a much larger effect on price when a car is in overall short supply than when it is not (-0.06 vs. -0.014), there is no difference for below median levels of inventory. This specification also illuminates our previous finding that days to resupply only affected prices for low levels of inventory. We had hypothesized that this is the case because a dealer with low inventory is more likely to run out of cars to sell when that dealer has to wait longer for new supply. This suggests that the price for “hot” cars, i.e. cars that are in overall short supply and for which dealers are less likely to have high inventory levels would also be sensitive to when the next shipment arrives. Consistent with this, the interaction between *DaysToResupply* and the “hot” indicator is positive and (marginally) significant.

7.2 Does inventory matter differently towards the end of the calendar month?

Next we examine whether the price-inventory relationship differs at the end of the month. As one can see across all specifications in the paper, cars that are sold at the end of the month sell for between 0.22% and 0.32% less than those sold at other times. This is because manufacturers and dealers impose non-linear sales targets on sales personnel which increases their incentive to sell additional cars at the end of the month. Due to these incentives we expect to observe

that at the end of the month dealers are less likely to hold out for a high price when inventory levels are low than they would at other times.

To test whether this is the case we replicate our basic specification using interactions of our inventory variables and days to resupply variables with the *EndOfMonth* indicator. The results in column 3 of Table 5 are not consistent with this hypothesis: while the *EndOfMonth* indicator continues to show that prices are lower at the end of the month, the price-inventory relationship seems not to be any different during that time.

7.3 Does inventory matter differently when a dealer's lot is full?

Finally, we expect to find a stronger price effect of inventory when a dealer lot is close to maximum capacity. This is because the inventory cost associated with the last few cars which the dealer can store on his lot may be particularly high. For example, the dealer might have to use customer parking, thereby decreasing revenue, or reduce spacing between parked cars, increasing the danger of damaging cars on the lot. Hence, we expect that inventory that is close to lot capacity will decrease the dealer's opportunity cost of selling a car and thus decrease transaction prices.

To test this hypothesis we add to the specification the indicator *LotFull* that is one when the total inventory held by the dealer across all cars is 95% or more of his maximum inventory over the sample period. We find no support for the hypothesis: the estimated coefficient on *LotFull* is insignificantly different from zero (see column 2 of Table 5). One possible explanation for this finding is that dealers are not normally size constrained. Therefore, reaching 95% of their maximum inventory over the sample period may not a good indicator of sharply increased inventory cost.

8 Inventory vs. price discrimination effects

We would like to get some sense of how much of the variation in car prices can be explained with price discrimination through bargaining and how much can be explained with inventory fluctuations. This is important because negotiated prices in car retailing are usually attributed to the fact that dealers try to price discriminate among consumers. Since we measure only a subset of what dealers observe when price discriminating among consumers, we cannot estimate precisely what fraction of the price variation can be attributed to these two factors. We can, however, determine how this fraction changes as a function of the type of dealership.

We measure and compare the average effect on price due to demographics and inventory using two indices. One is an index of the component of predicted price attributable to demo-

graphic factors,¹⁹ the other the component of predicted price attributable to inventory factors. We calculate these two indices as follows: We estimate a full specification (Table 5, column 1 with the addition of *LocalInventory*, *SubstituteInventory*, and *LotFull*) on a sample excluding no-haggle dealerships (see below). From the vector of estimated coefficients, $\hat{\beta}$, we extract two subvectors, $\hat{\beta}_D$ and $\hat{\beta}_I$, which are the vectors of coefficients for the demographic covariates and inventory covariates, respectively. The two indices are the products of these two coefficient subvectors and their corresponding data submatrices. Using the notation of equation (6), the demographic index is $D\hat{\beta}_D$ where D includes income, race, home ownership, and all other consumer demographics. The inventory, or supply-side index is $I\hat{\beta}_I$, where I includes the inventory and days to resupply measures, whether the car was a dealer trade, and *LocalInventory*, *SubstituteInventory*, and *LotFull*. Note that neither index includes the portion of the predicted price attributable to car and transaction characteristics (car fixed effects, vehicle cost, model recency, competition, weekend, region, and month). We measure the contribution of the two sets of factors to the overall variation in negotiated prices by comparing the relative magnitudes of the two indices. For each observation in the dataset, we divide the absolute value of the inventory index by the sum of the absolute values of both indices. Intuitively, this ratio measures the movement in price due to inventory versus that due to demographics for each observation when also controlling for other covariates. Averaging this ratio across observations, we find that for “haggle” dealerships the inventory measures explain, on average, 49.1% of the combined inventory and demographic components of the predicted price.

We can now test our empirical prediction about the relative magnitude of the inventory and price discrimination effects as a function of the type of dealership. Since many consumers have a high disutility of bargaining, some dealerships promise consumers a “no-haggle” price. “No haggle” does not mean fixed prices. These dealerships continue to set car prices daily based on their inventory and demand conditions; the no-haggle policy means that salespeople are discouraged from varying price across buyers who arrive on a particular day. Consequently, we expect that inventory fluctuations explain a larger percentage of price variation at “no-haggle” dealerships. We test this prediction for AutoNation, the largest chain of car dealerships in the US. AutoNation owns close to 300 dealerships nationwide (representing many different makes) and in 1996 adopted a policy of not negotiating with consumers. Our final data contains 45,280 transactions at AutoNation dealerships (out of a total 351,916 transactions). We expect inventory fluctuations in these dealerships to play a bigger relative role than in traditional

¹⁹We know from prior work that prices vary substantially with demographics, even after controlling for dealer fixed effects (Scott Morton, Zettelmeyer, and Silva-Risso 2003). This indicates that our demographics can serve as (imperfect) measures of the price discrimination.

dealerships.²⁰

We run separate regressions for transactions at AutoNation dealerships. We again calculate the indices described above. We find that the proportion of the inventory and demographic related price variation explained by inventory measures is 73.9% in AutoNation dealerships compared to 49.1% in the haggle dealerships.

These calculations confirm that inventory plays a larger role relative to demographics in no-haggle dealerships. On the face of it, these findings seem to indicate that salespeople at AutoNation dealerships still condition prices on consumer demographics. However, recall that the empirical correlation between transaction prices and consumer demographics could be driven by differing search methods or by outside opportunities that vary systematically by demographic characteristics, rather than price discrimination. While we are unable to distinguish between these two reasons for demographics and price to be correlated, we are able to show that inventory plays a larger role relative to demographics in the no-haggle setting.

9 Conclusion

We conclude from the evidence presented in the paper that local dealer inventory has a statistically and economically significant effect on the prices at which new cars are sold in the United States. A dealership moving from a situation of shortage to an average inventory level lowers transaction prices by about 1% *ceteris paribus*, corresponding to 15% of average dealer margins or \$250 on the average car. Additionally, shorter resupply times also decrease transaction prices for cars in high demand. We present some extensions of our results showing that local inventory also affects transaction prices and also that the effect of inventory is stronger for car models in high demand.

These results suggest that there is a valid economic, non-price-discriminatory reason for auto dealers to avoid posting prices for their cars. Since demand and supply conditions change constantly, the opportunity cost of using fixed, posted prices is substantial. Of course, setting price by negotiation opens the door to price discrimination. Interestingly, the dealerships that are trying to appeal to consumers who don't like price discrimination choose no-haggle pricing, not posted prices, as a method of sale.²¹ These dealerships, like traditional dealerships, appear to find the menu costs of maintaining posted prices that accurately reflect vehicle scarcity, or the opportunity costs of using posted prices that do not, too high to justify such a policy. We find

²⁰Saturn dealerships also have a policy of not haggling with customers. However, we have such few Saturn observations in our dataset that we cannot estimate all inventory and demographic effects on a Saturn-only sample.

²¹carsdirect.com is the only retailer we know of to set posted prices for cars. Note that its large scale reduces the "per sale" menu costs.

that price variation at no-haggle dealerships is less well predicted by consumer demographics, while inventory considerations remain important. The share of the predicted price attributable to either inventory or demographics that is due to inventory is 49% in traditional dealerships. For so-called “no-haggle” dealerships, inventory explains 74% of the combined inventory and demographic components of the predicted price. This is consistent with salespeople at no-haggle dealerships engaging in less price discrimination.

We conclude that price differences in car retailing are in part the result of scarcity rents, and function partially to efficiently allocate particular cars that are in restricted supply to those customers who value them most highly.

References

- 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.
- CARLTON, D. (1978): “Market Behavior with Demand Uncertainty and Price Inflexibility,” *American Economic Review*, 68, 571–587.
- COPELAND, A., W. DUNN, AND G. HALL (2005): “Prices, Production, and Inventories over the Automotive Model Year,” Working paper, Yale University.
- COPELAND, A., AND G. HALL (2005): “The Response of Prices, Sales, and Output to Temporary Changes in Demand,” Discussion paper, Yale University.
- DANA, J. (2001): “Competition in price and availability when availability is unobservable,” *RAND Journal of Economics*, 32(4), 497–513.
- GALLEGO, G., AND G. V. RYZIN (1994): “Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons,” *Management Science*, 40(8), 999–1020.
- HALL, G., AND J. RUST (2000): “An Empirical Model of Inventory Investment by Durable Commodity Intermediaries,” *Carnegie-Rochester Conference Series on Public Policy*.
- 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.
- 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	25,757.24	10,361.88	5213	126,490	351,916
Inventory	15.47	14.87	1	80	351,916
DaysToResupply	8.44	12.22	1	147	351,916
#Cars arriving in week+1	3	4.23	0	71	351,916
#Cars arriving in week+2	2.92	4.17	0	88	351,916
#Cars arriving in week+3	2.84	4.16	0	76	351,916
#Cars arriving in week+4	2.75	4.12	0	79	351,916
TradedCar	0.07	0.26	0	1	351,916
# Competing Dealers	4.31	3.15	0	23	351903
%Black	4.2	9.67	0	100	351,916
%Hispanic	15.16	11.69	0	55.33	351,916
HispanicName	0.15	0.35	0	1	351,916
%Asian	10.95	12.14	0	100	351,916
Asian	0.07	0.26	0	1	351,916
Female	0.29	0.45	0	1	351,916
Income	61,383.28	26,187.54	11,318	150,000	351,916
Income ²	4.45E+09	3.95E+09	1.28E+08	2.25E+10	351,916
%CollegeGrad	32.8	17.4	0	100	351,916
%LessHighSchool	12.36	11.98	0	100	351,916
%HouseOwnership	66.69	24.79	0.14	100	351,916
%Professional	16.95	8.71	0	100	351,916
%Executives	18.17	8.22	0	100	351,916
%BlueCollar	25.27	15.33	0	100	351,916
%Technicians	3.07	2.09	0	100	351,916
MedianHouseValue	240,999.25	113,620.25	7500	500,000	351,916
CustomerAge	42.36	13.05	16	102	351,916
Age> 64	0.06	0.24	0	1	351,916
VehicleCost	0	0.05	-0.64	0.88	351,916
Model Age 5-13 Months	0.68	0.47	0	1	351,916
Model Age > 14 Months	0.05	0.23	0	1	351,916
Weekend	0.37	0.48	0	1	351,916
EndOfMonth	0.22	0.41	0	1	351,916
EndOfYear	0.02	0.16	0	1	351,916

Table 2: Price effects of inventory[†]

Dep. Var. ln(price)	Fixed Effects		
	Car, Dealer, Month	Car*Dealer, Month	Car*Dealer, Month*Segment*DMA
Inventory (1-10)	-.085** (.0044)	-.092** (.005)	-.064** (.0051)
Inventory (11+)	-.018** (.0013)	-.016** (.0018)	-.013** (.0019)
DaysToResupply (1-4)	-.015 (.0095)	-.0052 (.0095)	-.0035 (.0095)
DaysToResupply (4+)	.0005 (.001)	.0013 (.0011)	.0011 (.0011)
TradedCar	-.0075 (.035)	-.0072 (.035)	-.0028 (.035)
EndOfYear	.017 (.077)	-.0033 (.076)	-.09 (.076)
EndOfMonth	-.31** (.025)	-.28** (.025)	-.38** (.025)
Weekend	.07** (.022)	.066** (.022)	.067** (.022)
VehicleCost	.98** (.22)	.99** (.23)	.99** (.23)
Model Age 5-13 Months	-.98** (.033)	-.94** (.034)	.048 (.043)
Model Age > 14 Months	-1.9** (.076)	.54** (.076)	.38** (.096)
Income	-.00001** (2.3e-06)	-8.7e-06** (2.3e-06)	-8.5e-06** (2.3e-06)
Income ²	7.3e-11** (1.2e-11)	6.2e-11** (1.2e-11)	6.1e-11** (1.2e-11)
%CollegeGrad	-.0018 (.0015)	-.0013 (.0015)	-.0016 (.0015)
%LessHighSchool	.0099** (.002)	.0095** (.002)	.0096** (.002)
%HouseOwnership	-.00027 (.00068)	-.00039 (.00068)	-.00045 (.00067)
MedianHouseValue	-3.8e-07* (1.8e-07)	-3.9e-07* (1.7e-07)	-4.1e-07* (1.7e-07)
%Professional	-.0029 (.0022)	-.0033 (.0022)	-.0033 (.0022)
%Executives	.0018 (.0023)	.00085 (.0022)	.0012 (.0022)
%BlueCollar	-.0012 (.0018)	-.0012 (.0017)	-.00096 (.0017)
%Technicians	-.011* (.0051)	-.013* (.005)	-.013** (.005)
CustomerAge	.014** (.00096)	.014** (.00095)	.014** (.00095)
Age > 64	.0056 (.051)	-.0099 (.051)	.0018 (.051)
Female	.14** (.022)	.14** (.022)	.14** (.022)
%Asian	-.0024* (.00097)	-.002* (.00096)	-.002* (.00095)
%Black	.012** (.0012)	.011** (.0012)	.011** (.0012)
%Hispanic	.0056** (.0017)	.0051** (.0017)	.0049** (.0017)
Constant	1008** (.49)	1008** (.16)	1009** (.28)
Observations	351916	351916	351916
Adj. R-squared	0.974	0.975	0.976

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

[†] Unreported are car, dealer, and monthly fixed effect (column 1); car*dealer and monthly fixed effect (column 2); car*dealer and month*segment*DMA fixed effect (column 3).

All coefficients are multiplied by 100.

Table 3: Robustness checks[†]

Dep. Var. ln(price)	(1) DTR with # of arrivals	(2) DTR at low inventory	(3) Local Inventory	(4) Substitute Cars	(5) Coarser Inven. Def.
Inventory (1-10)	-.092** (.005)	-.088** (.0053)	-.08** (.0054)	-.085** (.0056)	
Inventory (11+)	-.016** (.0019)	-.016** (.0018)	-.01** (.0019)	-.015** (.0019)	
DaysToResupply	.001 (.00099)	.00046 (.0011)	.00049 (.0011)	.00061 (.0011)	
DTR* LowInventory		.0035+ (.0019)	.0036+ (.0019)	.0045* (.002)	
#Cars arriving in week+1	-.00055 (.0036)				
#Cars arriving in week+2	-.0002 (.0036)				
#Cars arriving in week+3	.0019 (.0036)				
#Cars arriving in week+4	-.00083 (.0037)				
LocalInventory			-.0049** (.00039)		
Substitute- Inventory				-.0055** (.0013)	
Coarse Inventory (≤ Median)					-.069** (.0036)
Coarse Inventory (> Median)					-.0083** (.0011)
DTR (for Coarse Inventory)					.0028 (.027)
DTR (Coarse)* LowInventory					.048 (.066)
TradedCar	-.0073 (.035)	-.0074 (.035)	-.0073 (.035)	-.0045 (.037)	.004 (.036)
EndOfYear	-.0043 (.077)	-.0036 (.076)	-.004 (.076)	-.027 (.08)	-.025 (.08)
EndOfMonth	-.28** (.025)	-.28** (.025)	-.28** (.025)	-.28** (.026)	-.28** (.026)
Weekend	.063** (.021)	.063** (.021)	.066** (.021)	.058** (.022)	.059** (.022)
VehicleCost	.99** (.23)	.99** (.23)	.99** (.23)	.99** (.24)	.99** (.24)
Model Age 5-13 Months	-.94** (.034)	-.94** (.034)	-.88** (.034)	-.95** (.035)	-.93** (.035)
Model Age > 14 Months	-1.8** (.076)	-1.8** (.076)	-1.7** (.077)	-1.8** (.078)	-1.7** (.078)
Constant	1008** (.16)	1008** (.16)	1008** (.16)	1008** (.17)	1009** (.17)
Observations	351916	351916	351916	332020	332020
Adj. R-squared	0.975	0.975	0.975	0.976	0.976

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

[†] Unreported are car*dealer, monthly fixed effect, and the demographic variables reported in Table 2.

All coefficients are multiplied by 100.

Table 4: Instrumental variables[†]

Dep. Var. ln(price)	(1) OLS	(2) IV
Inventory (1-10)	-.11** (.014)	-.17** (.063)
Inventory (11+)	-.013** (.0038)	-.012 (.015)
DaysToResupply	.00032 (.0021)	-.015 (.0095)
DTR*LowInventory	.0044 (.0042)	.052** (.019)
TradedCar	.11 (.094)	.05 (.11)
EndOfYear	-.17 (.19)	-.14 (.19)
EndOfMonth	-.32** (.062)	-.32** (.062)
Weekend	.01 (.053)	.024 (.052)
VehicleCost	101** (.51)	101** (.44)
Model Age 5-13 Months	-1.2** (.088)	-1.1** (.1)
Model Age > 14 Monthss	-2.2** (.17)	-2** (.19)
Constant	1013** (.43)	1013** (.64)
Observations	75213	75213
Adj. R-squared	0.945	

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

[†] Unreported are car*dealer, monthly fixed effect, and the demographic variables reported in Table 2.

All coefficients are multiplied by 100.

Table 5: Additional specifications[†]

Dep. Var. ln(price)	(1) Hot Cars	(2) End of Month	(3) Lot Full
Inventory (1-10)	-.089** (.0058)	-.086** (.0057)	-.088** (.0053)
Inv.(1-10)*HotCar	-.0015 (.011)		
Inv.(1-10)*EndOfMonth		-.01 (.0093)	
Inventory (11+)	-.014** (.0019)	-.016** (.0019)	-.016** (.0018)
Inv.(11+)*HotCar	-.046** (.0078)		
Inv.(11+)*EndOfMonth		-.00027 (.0025)	
DaysToResupply	.00053 (.001)	.00033 (.0012)	.00046 (.0011)
DTR*LowInventory		.0024 (.0021)	.0035+ (.0019)
DTR*HotCar	.0046+ (.0028)		
DTR*EndOfMonth		.00062 (.0023)	
DTR*LowInv*EndOfMonth		.0052 (.0044)	
Inv>95%MaxInv			-.0024 (.078)
TradedCar	-.0066 (.035)	-.0072 (.035)	-.0074 (.035)
EndOfYear	-.0048 (.076)	-.0028 (.076)	-.0036 (.076)
EndOfMonth	-.28** (.025)	-.22** (.068)	-.28** (.025)
Weekend	.064** (.021)	.063** (.021)	.063** (.021)
VehicleCost	.99** (.23)	.99** (.23)	.99** (.23)
Model Age 5-13 Months	-.94** (.034)	-.94** (.034)	-.94** (.034)
Model Age > 14 Months	-1.8** (.076)	-1.8** (.076)	-1.8** (.076)
Constant	1008** (.16)	1008** (.16)	1008** (.16)
Observations	351916	351916	351916
Adj. R-squared	0.975	0.975	0.975

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

[†] Unreported are car*dealer, monthly fixed effect, and the demographic variables reported in Table 2.

All coefficients are multiplied by 100.

Figure 1: Example of inventory movement for 3 cars at one Honda dealer in Jan/Feb 1999

Franchise 82 (Honda)

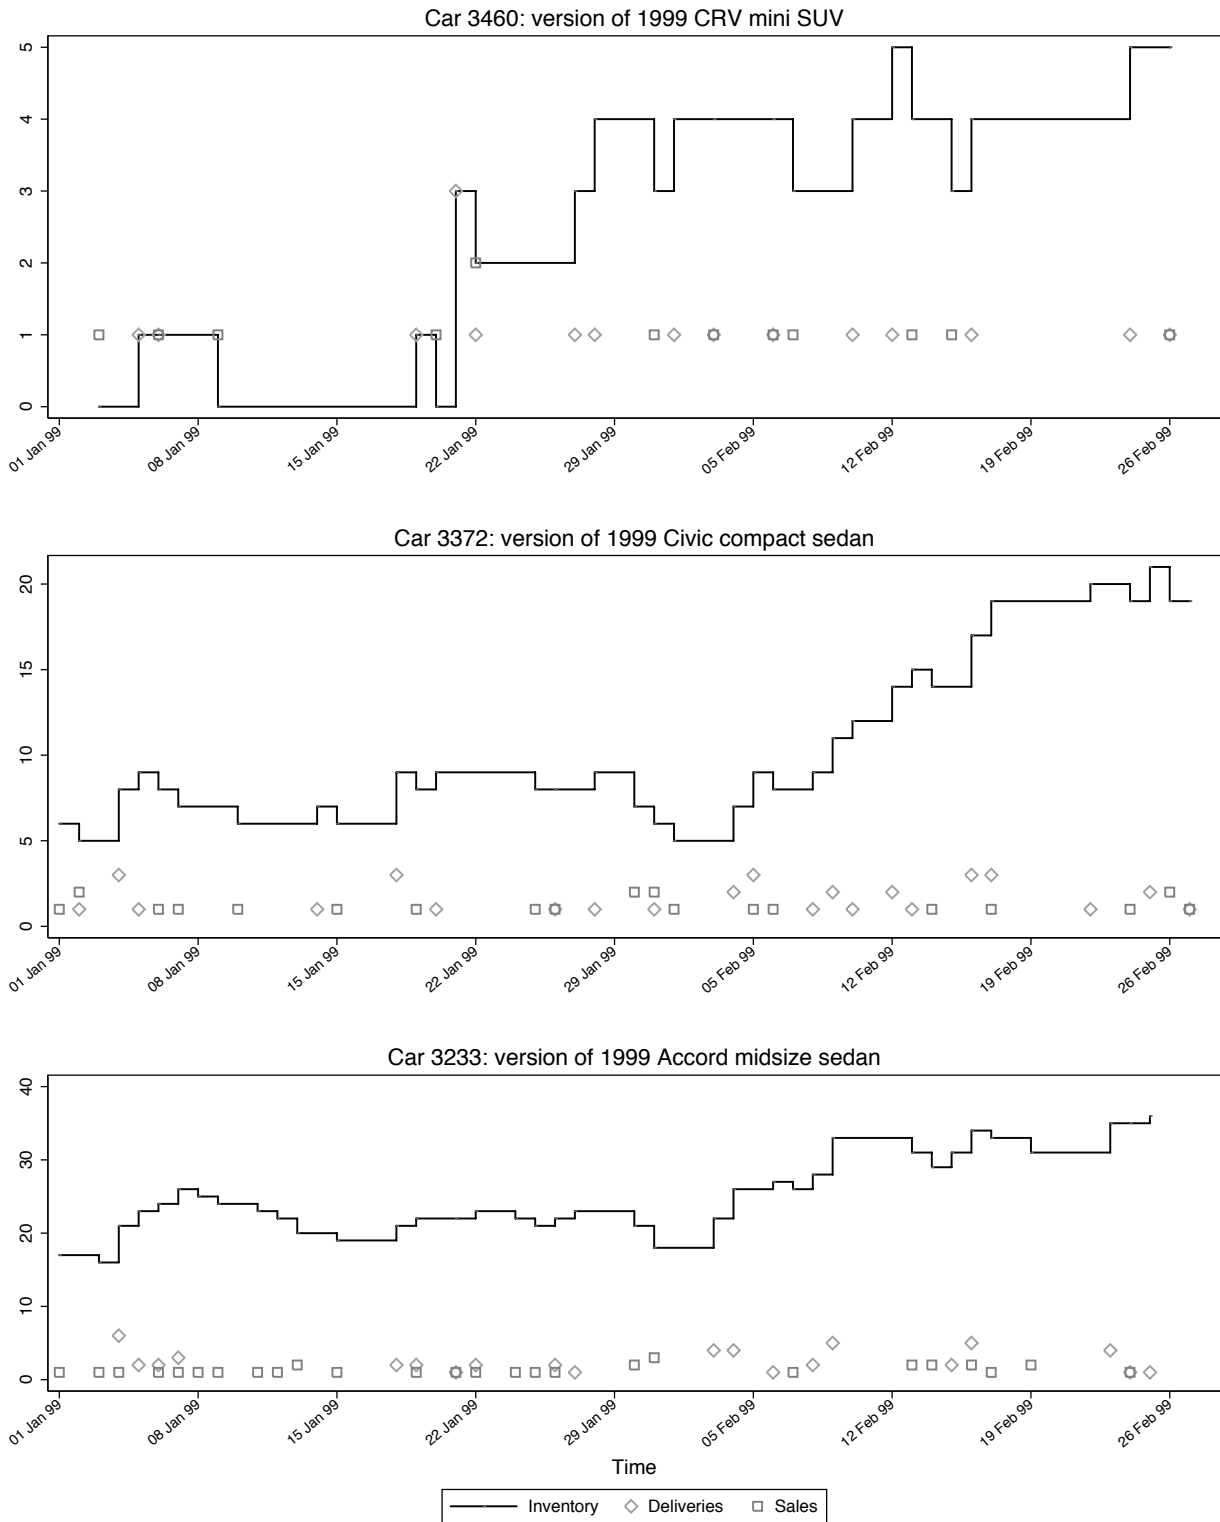


Figure 2: Example of inventory movement for 3 cars at one Chevrolet dealer in Jan/Feb 1999

Franchise 2193 (Chevrolet)

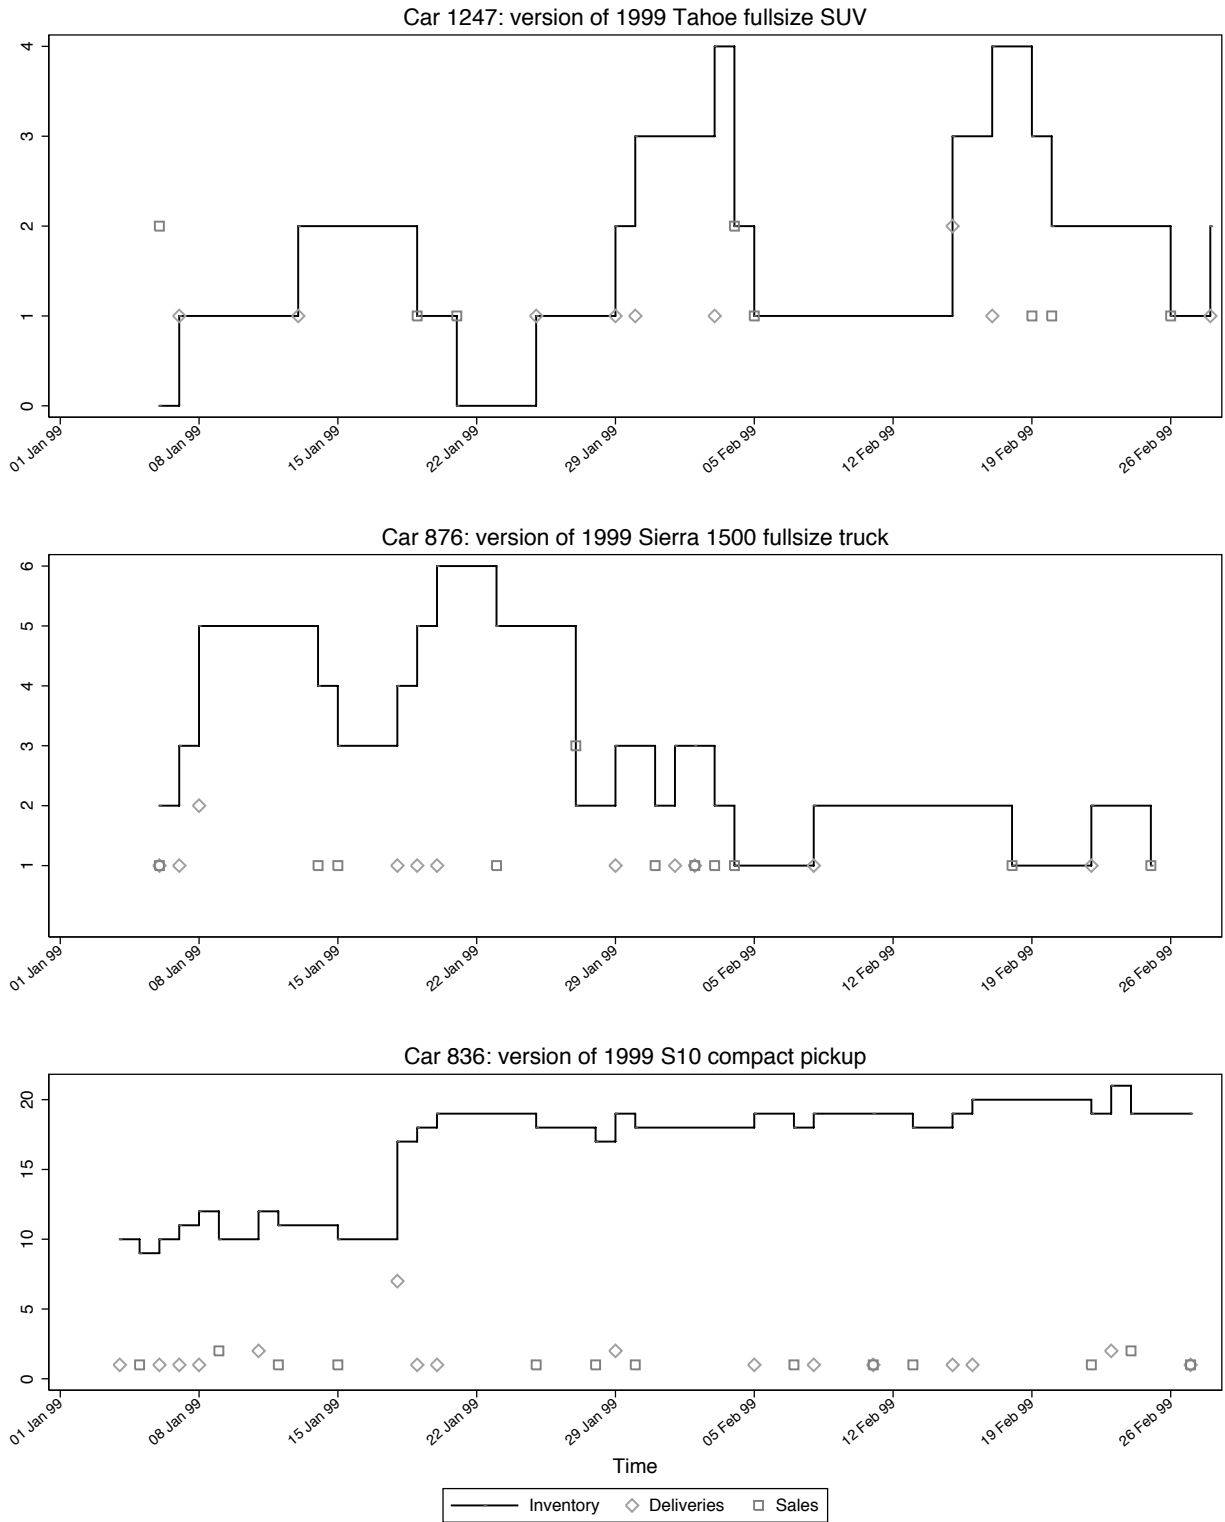


Figure 3: Example of inventory movement for 3 cars at one Mercedes dealer in Jan/Feb 1999

Franchise 701 (Mercedes)

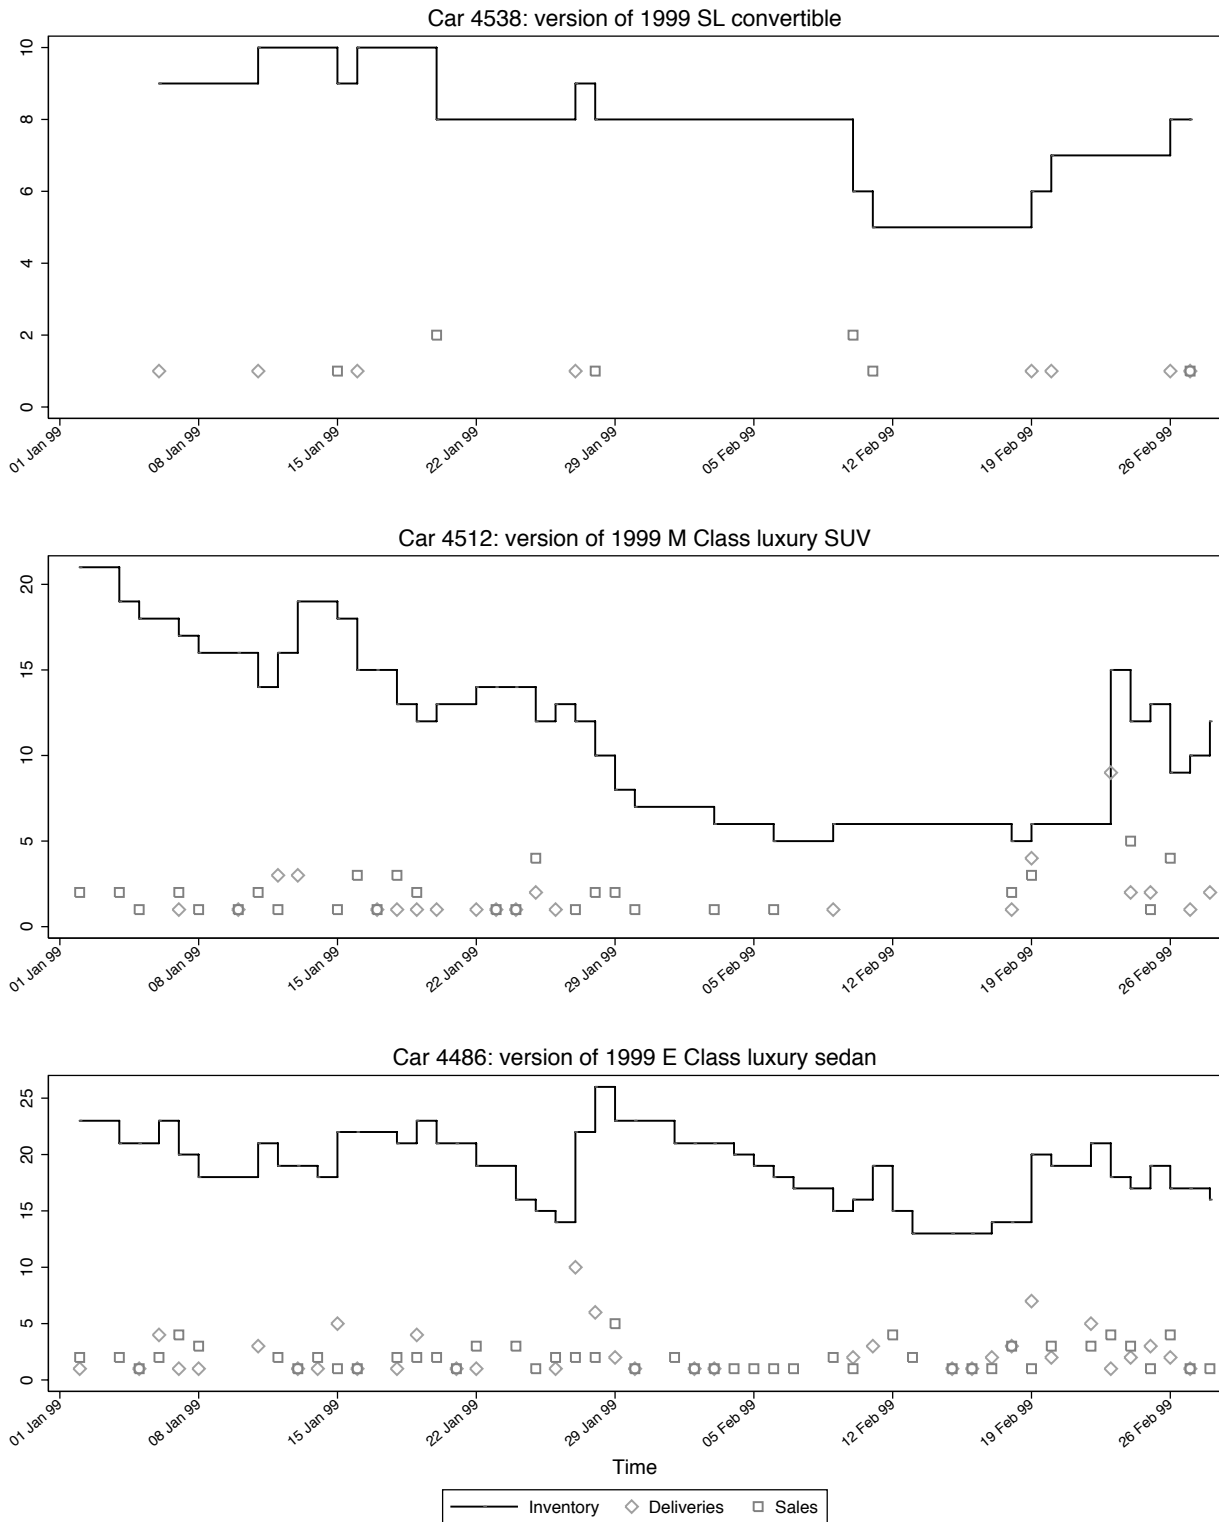


Figure 4: Distribution of daily inventories (at the inventory “car” level)

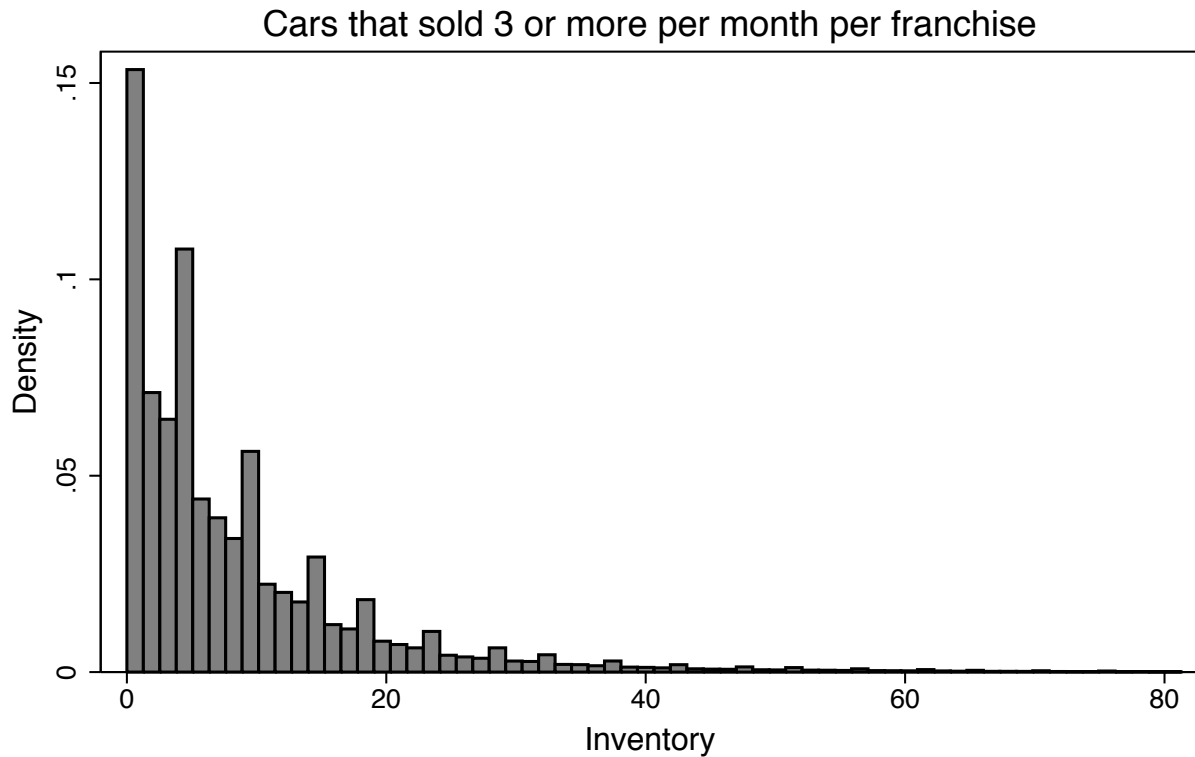
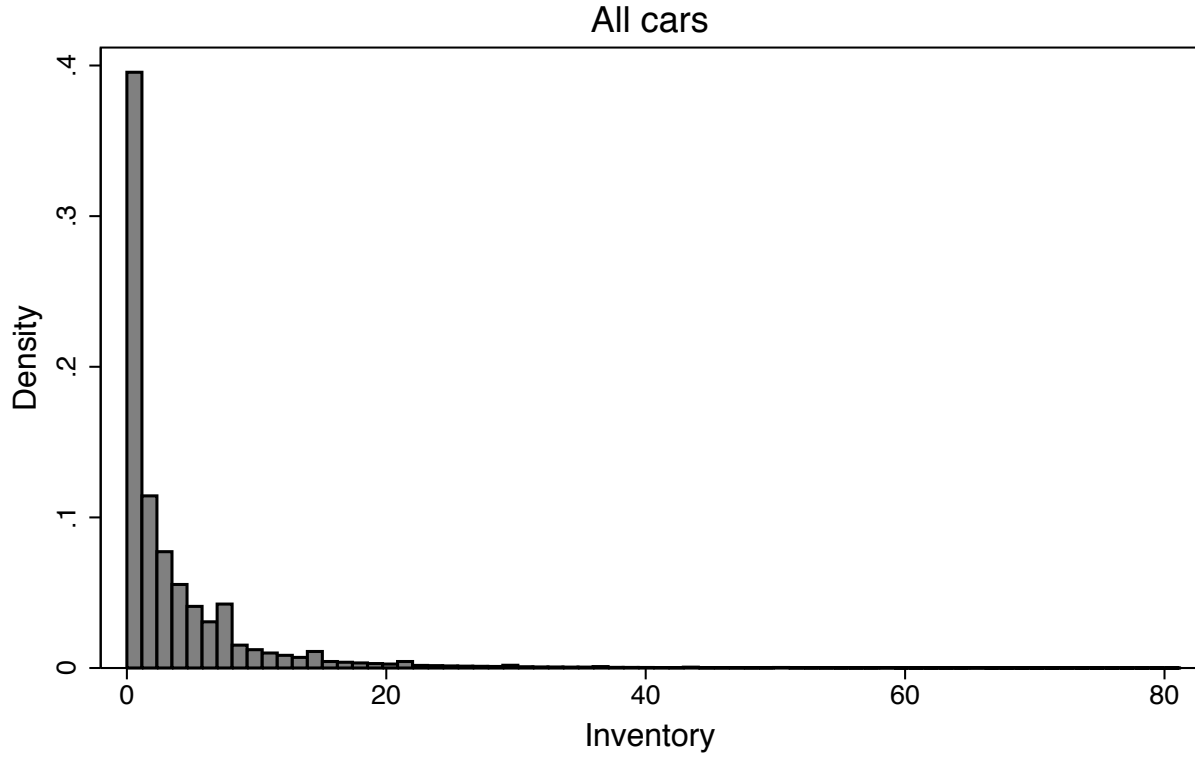


Figure 5: Distribution of daily 'days to resupply' (at the inventory "car" level)

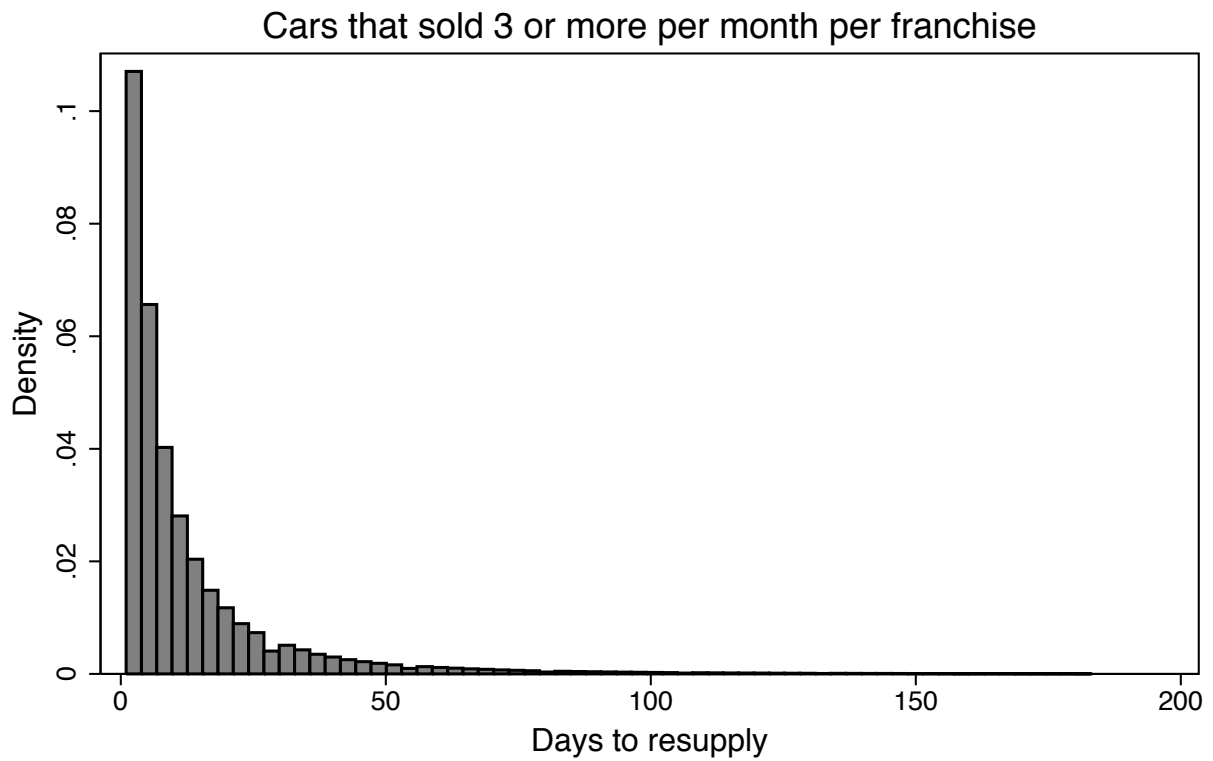
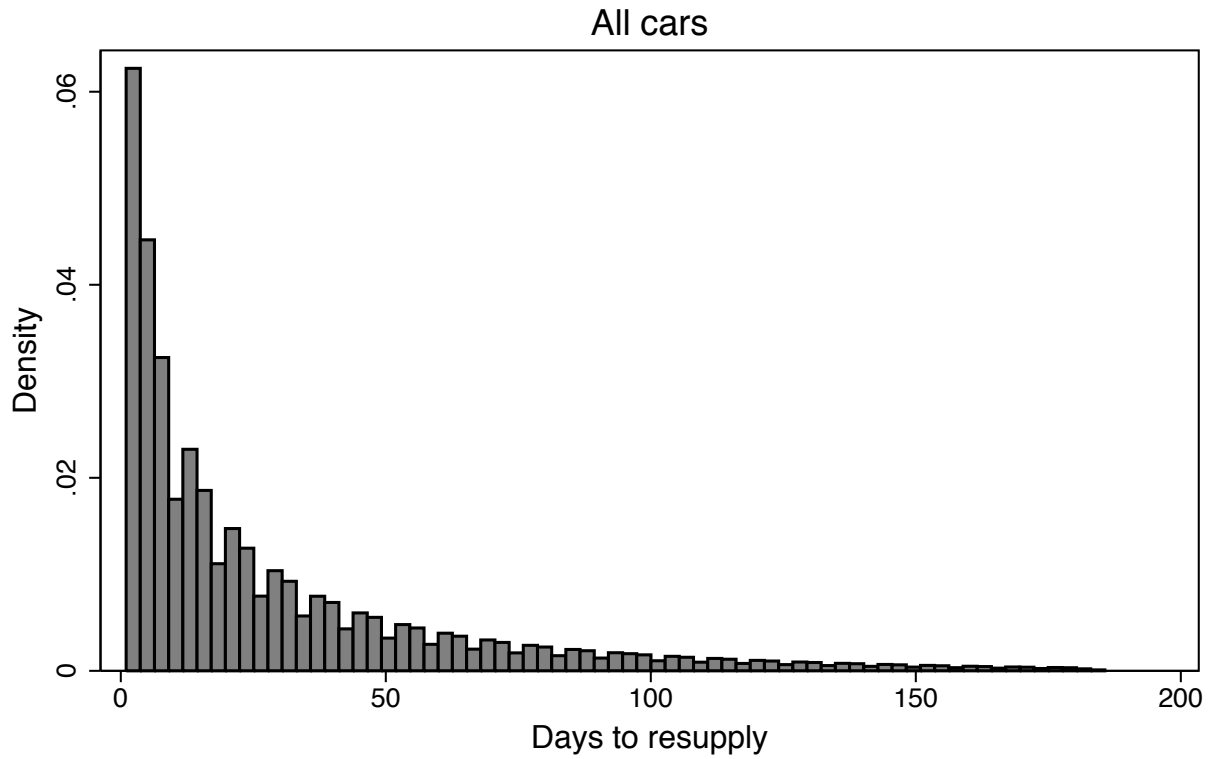


Figure 6: Distribution of the number of vehicles (at the inventory “car” level) arriving on a day (conditional on supply, i.e. on at least one car arriving)

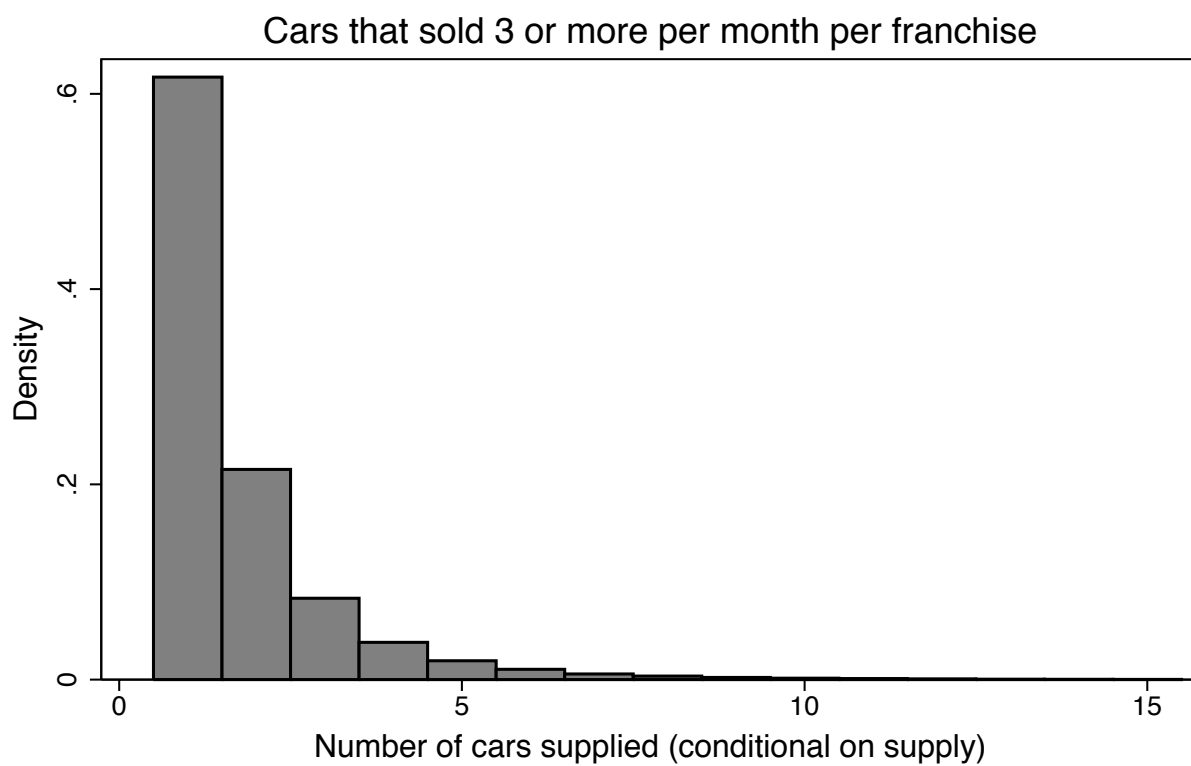
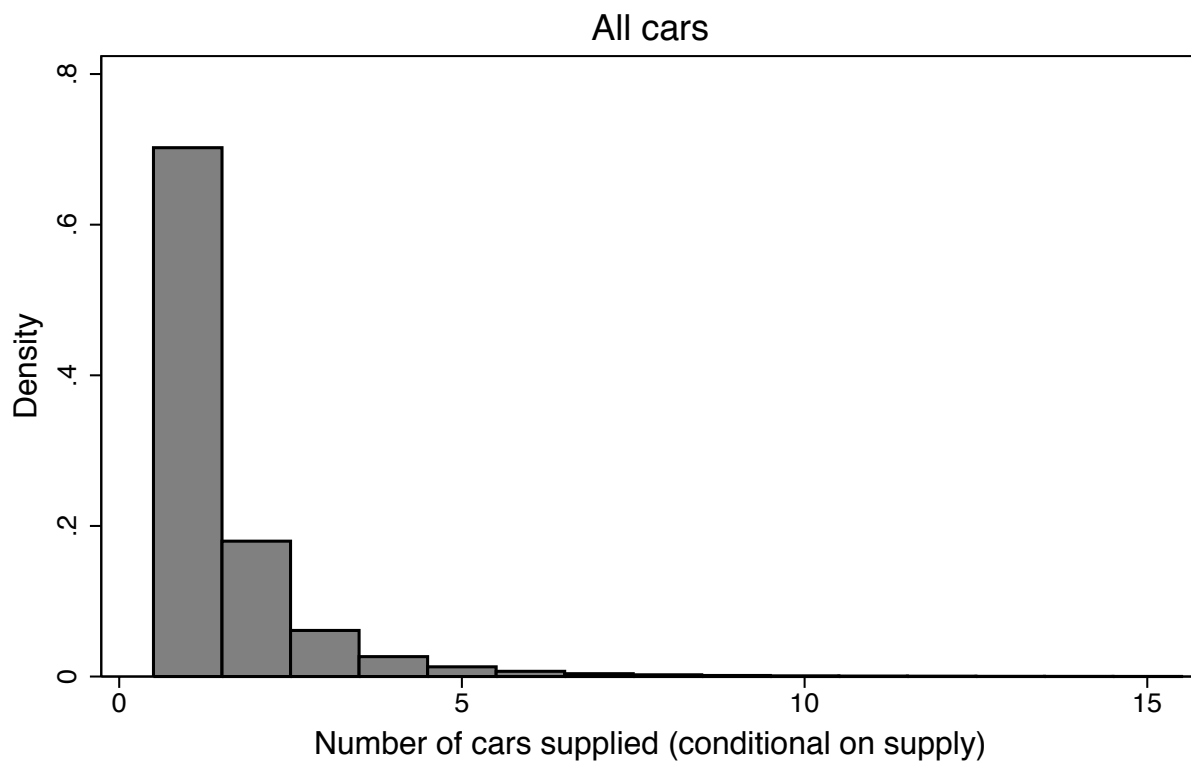


Figure 7: Comparison of the price-inventory relationship across car-dealer sales volume quartiles with the full sample results

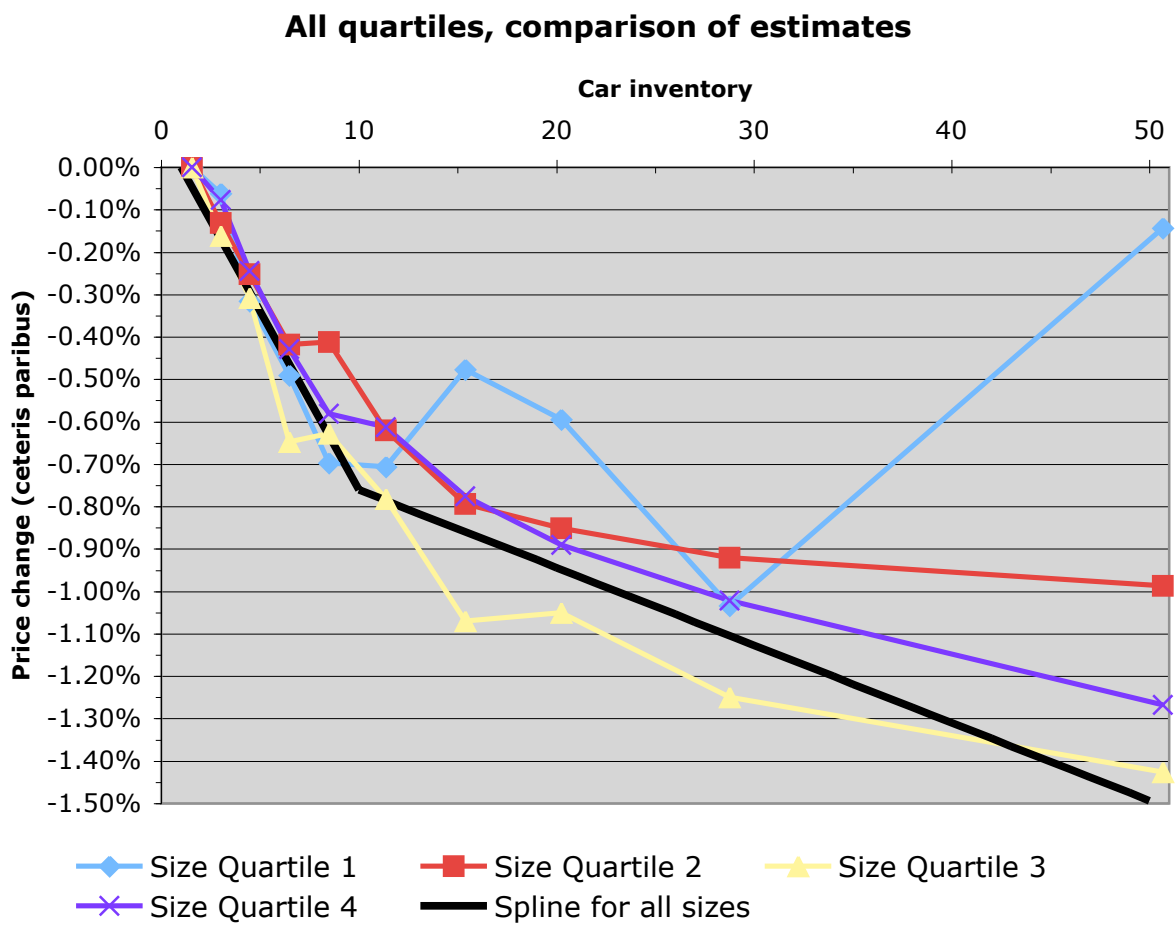


Figure 8: Comparison of the full sample estimates with the estimates and associated confidence intervals for each car-dealer sales volume quartile

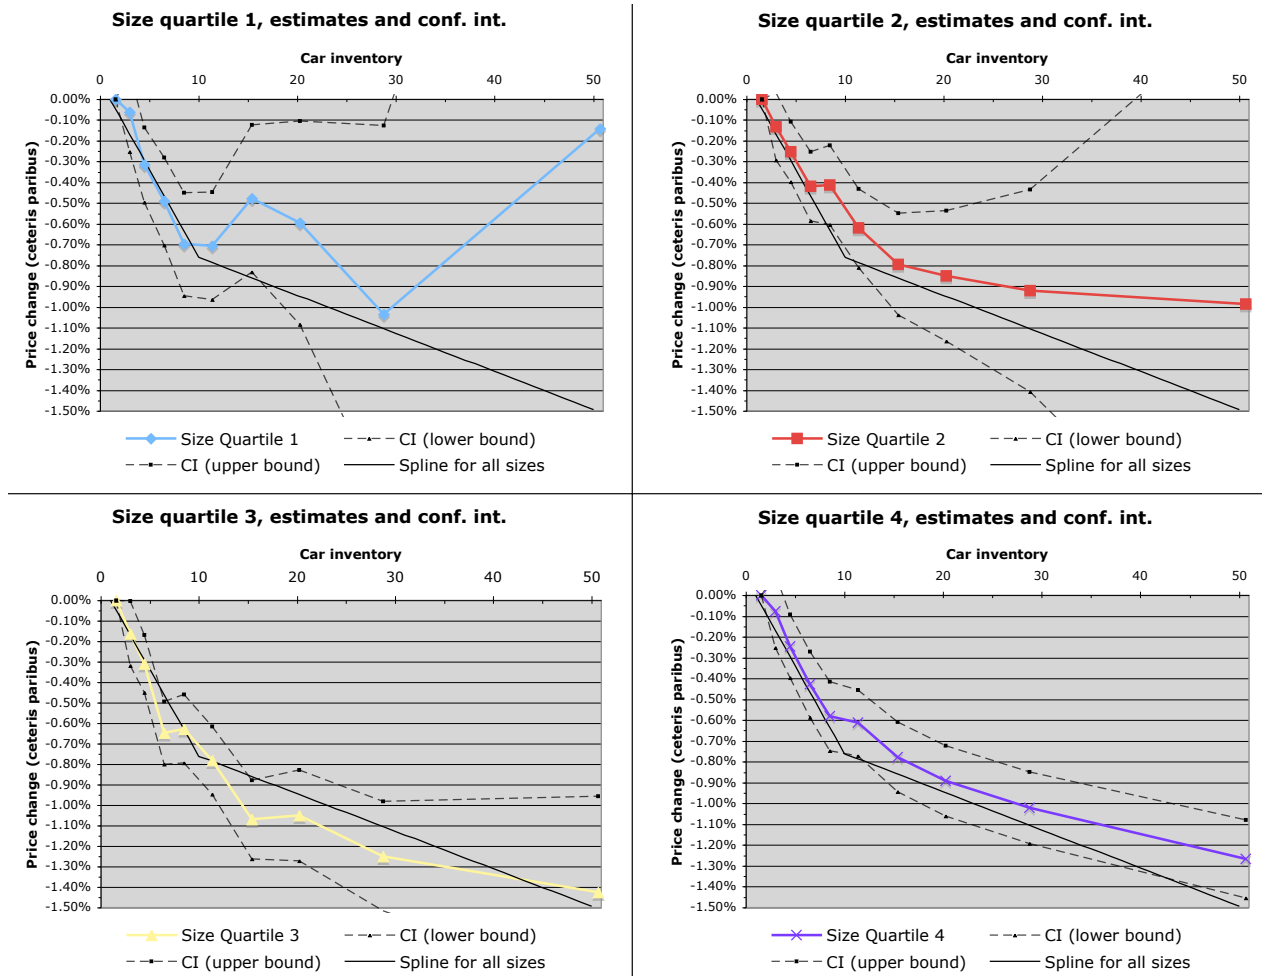


Figure 9: Number of observations by size quartile and inventory dummy for the estimations in Figures 7 and 8.[†]

Inventory Decile	Mean Inventory	Size Quartile				Total
		1	2	3	4	
1	1.6	7,814	9,679	8,747	6,931	33,171
2	3.0	4,305	5,720	5,838	5,381	21,244
3	4.5	7,390	10,771	11,659	12,158	41,978
4	6.5	5,654	8,550	10,293	12,471	36,968
5	8.5	3,815	6,376	8,910	11,907	31,008
6	11.4	4,237	8,523	12,758	21,524	47,042
7	15.4	1,958	4,152	8,355	18,818	33,283
8	20.3	1,144	2,639	6,599	23,699	34,081
9	28.8	346	1,285	4,741	30,453	36,825
10	50.6	37	164	1,436	34,679	36,316
	Total	36,700	57,859	79,336	178,021	351,916