

UC Berkeley

Econ 221 - Industrial Organization Seminar

Title

Inventory Fluctuations and Price Discrimination: The Determinants of Price Variation in Car Retailing

Permalink

<https://escholarship.org/uc/item/6pf5s593>

Authors

Zettelmeyer, Florian
Scott Morton, Fiona
Silva-Risso, Jorge

Publication Date

2005-08-30

Inventory Fluctuations and Price Discrimination: The Determinants of Price Variation in Car Retailing*

Florian Zettelmeyer

University of California at Berkeley and NBER

Fiona Scott Morton

Yale University and NBER

Jorge Silva-Risso

University of California at Los Angeles

July 2005

Draft – Preliminary

*We thank Severin Borenstein, Jose Silva, Candi Yano, and especially Meghan Busse and Thomas Hubbard for helpful ideas and comments. Addresses for correspondence: Haas School of Business, UC Berkeley, Berkeley CA 94720-1900; School of Management, Yale University, PO Box 208200, New Haven CT 06520-8200; Anderson School at UCLA, 110 Westwood Plaza, Los Angeles, CA 90095. E-mail: florian@haas.berkeley.edu, fiona.scottmorton@yale.edu, jorge.silva-risso@anderson.ucla.edu

Inventory Fluctuations and Price Discrimination: The Determinants of Price Variation in Car Retailing

Abstract

We show that variation in prices for new cars are in part the result of scarcity rents. Consistent with theoretical predictions, a dealership moving from a situation of shortage to one of ample inventory lowers transaction prices by about 1% *ceteris paribus*, or \$230 on the average car. Shorter resupply times also decrease transaction prices for cars in excess demand. Transaction prices for cars in excess supply are more responsive to inventory variables at the end of the model year. Consumers pay lower prices when dealers' total inventory is high and when inventory of the purchased car at neighboring dealers is relatively high. For traditional dealerships, inventory explains 21% of the combined inventory and demographic components of the predicted price. For so-called "no-haggle" dealerships, the percentage explained by inventory increases to 51% (for Autonation dealers) and 65% (for Saturn dealer).

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 fixed in the short term at the inventory on the 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 and demand into the offered price. As a result, the opportunity cost to the dealer of selling a car—and therefore 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, in 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 the amount of inventory of that specific car in his lot and the number of days remaining until a shipment of new inventory arrives. This 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, for example, Gallego and Ryzin (1994) and Papastavrou, Rajagopalan, and Kleywegt (1996)). Our example shows that it is possible for these results to carry over to a negotiated price environment. The intuition behind the result 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 resupply result, 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.

The empirical section of the paper provides evidence for both of these relationships. A dealership moving from a situation of shortage to one of ample inventory lowers transaction prices by about 1% *ceteris paribus*, or \$230 on the average car. Additionally, shorter resupply times also decrease transaction prices, as models of optimal pricing suggest they should. These effects vary according to the excess demand of the car model in question. Cars in substantial excess demand show stronger effects of inventory on price and of resupply times on price compared to cars in excess supply. Inventory factors have a strong effect on the prices of cars in excess supply at the end of the model year, however. Consumers pay lower prices when a dealer's lot is full (about 1% less) and when inventory of the purchased car at neighboring dealers is relatively high (0.4% less). Finally, we find that the share of the predicted price attributable to either inventory or demographics that is due to inventory is 21% in our sample. For so-called "no-haggle" dealerships, inventory explains 51% (for Autonation dealers) and 65% (for Saturn dealer) of the combined inventory and demographic components of the predicted price.

The paper proceeds as follows. In section 2 we develop a simple model to illustrate the relation of prices and inventory levels. We derive empirical predictions and discuss estimation issues. In section 3 we describe our data and discuss measurement issues in inventory. Section 4 presents the results and section 5 concludes the paper.

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 Example

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 f_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 shipments are determined at least 45 days before cars are delivered and typically 90 days before. Within that time period, dealers cannot obtain additional cars from the manufacturer, reduce their order, or alter its composition.¹ Because shipments are exogenous in the short run, the marginal cost of shipped cars does not matter for the dealer’s pricing problem. If the dealer has no cars on the lot, he (obviously) cannot sell any cars until the next shipment. We assume that consumers drop out of the market if they find no inventory (or purchase from another dealer). 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.²

We assume that price is determined according to a standard Roth-Nash bargaining model. The price paid for a car by the consumer who arrives at the dealership at time 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 and λ with t .

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

This formula 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

¹However, they can trade vehicles with other dealers. We do not consider this possibility in the model, but in the empirical analysis we control for inter-dealer trades.

²The strict lot size constraint has the same effect in this model as an inventory holding cost.

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.

The problem now consists of determining 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 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 $1 \leq t < T$ 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 this expression notice that the dealer will sell a car if the reservation price of the buyer exceeds the dealer's opportunity cost. The dealer obtains revenue of $\lambda(r_t - o_t) + o_t$ and enters the next period with $n - 1$ cars. If there are no gains from trade ($r_t < o_t$) the dealer sells no car and enters the next period with n cars. This 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 with the smaller of $n - 1 + S$ cars on his lot or a full lot, L . If $n - 1 + S > L$ the dealer needs to return $n - 1 + S - L$ cars to the manufacturer at return fee f per returned car. If the dealer does not sell a car in the last period of the inventory cycle he enters the next inventory cycle with $\min\{n + S, L\}$ cars and needs to pay a return cost on $\max\{0, n + S - L\}$ cars. Formally,

$$\begin{aligned} \Pi(n, T, c) = & Pr(r_T \geq o_T) (E_\lambda [E_r[\lambda(r_T - o_T) + o_T \mid r \geq o_T]] + \\ & \Pi(\min\{n - 1 + S, L\}, 1, c + 1) - f \max\{0, n - 1 + S - L\}) + \\ & Pr(r_T < o_T) (\Pi(\min\{n + S, L\}, 1, c + 1) - f \max\{0, n + S - L\}) \end{aligned} \quad (3)$$

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

$$\Pi(0, t, c) = \Pi(S, 1, c + 1) \quad (4)$$

Using (2), (3), and (4), 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		
	$t = 1$	$t = 2$	$t = 3$
$n = 1$	0.72	0.62	0.50
$n = 2$	0.33	0.24	0.15
$n = 3$	0.09	0.04	-0.003

To get an intuition for the dealer’s opportunity cost, 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 three opportunities to sell that car before he receives a replacement car and thus holds out for a high valuation buyer by setting the minimum offer it is willing to accept at 0.72. In the next period, the dealer has two opportunities to sell that vehicle before the next shipment, resulting in a lower opportunity cost for the vehicle. 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. Next, fix the number of days until the next shipment, for example $t = 3$, and consider the change in price as the dealer has more cars on the lot. If the dealer has only one car in inventory, he holds out for a higher valuation buyer than if he has two cars in inventory 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 the dealer’s opportunity cost can be 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 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):

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. These comparative static predictions are not unique to this setup but are shared across a class of models in operations

	Expected negotiated prices		
	$t = 1$	$t = 2$	$t = 3$
$n = 1$	0.79	0.72	0.62
$n = 2$	0.50	0.43	0.36
$n = 3$	0.31	0.28	0.25

research (see Yano and Gilbert (2002) for a review). These are models 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. One of the closest papers to our own in this line of research is a model by Gallego and Ryzin (1994) which characterizes the profit maximizing prices of a monopolist over a finite horizon as a function of the inventory and the time remaining until the deadline. Their model allows for a salvage value at the end of the (single) inventory cycle and is thus a good representation of pricing within one cycle in our model, except for the fact 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.

Our setup varies from this literature in two ways: we assume that prices are negotiated instead of being set by a monopolist, and we consider recurring shipments over an infinite horizon. However, the model intuitions are similar enough that we are reasonably certain that the comparative statics predictions from our example are robust to general lot size constraints, length of inventory cycles, and shipment quantities.³

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,

³We have also solved the model for many different parameters.

except at full capacity.⁴ 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	0.20	0.47
$n = 1$	0.72	0.70	0.51
$n = 2$	0.27	0.09	0.02
$n = 3$	0.003	0.0003	9E-06

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 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 days to next shipment are not dependent on whether shipment quantities and/or inventory cycles are chosen optimally by the dealer. We have anecdotal evidence that they are chosen by the manufacturer, and are not generally chosen in the best interest of the dealer. For example, manufacturers often force dealers to take delivery of low demand cars as a condition for obtaining some high demand cars.

2.2 Empirical predictions

We have several goals for the paper. First, we would like to see whether our two comparative statics hold:

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

Second, we expect that the inventory-price relationship will be stronger in some cases than others. Inventory should have a particularly large effect on price for transactions that take place near the end of the model year. The dealer's opportunity cost of selling these cars decreases as the arrival of the new model draws nearer, typically in September of any given year. This is because dealers expect a decrease in the consumer's willingness to pay for the old model since the resale value of a vehicle is tied to its model year and the market does not distinguish between cars produced at the beginning and at the end of the model year.

⁴Since the firm is resupplied in period 1, it never has 0 cars in period 1.

The inventory-price relationship may also vary by the popularity of cars. This is because there may be some car models for which the dealer can get extra supply on short notice if the demand realization is higher than expected. This happens through dealer trades, and can cause the inventory of cars with different net supply situations to have different relationships between inventory and price.

For example, a car in substantial excess supply (“dud” car) will be one that the dealer can obtain on short notice if there is an unexpected increase in the dealer’s demand. The general plentiful supply of these cars may reduce dealers’ propensity to adjust price to their own low inventory. In contrast, a dealer will be hard-pressed to obtain additional cars in short supply (“hot” cars) if the he runs out of inventory. In the empirical specification, we distinguish between “hot”, “dud”, and regular cars according to their average sale times across all dealers in California. We expect to see weaker inventory effects for dud cars than for “hot” and regular cars.

We also expect that the amount of inventory in the dealer’s local geographic area will affect whether dealer inventory affects car prices. This is because the ease with which a dealer can obtain a trade from another dealer depends on the number of cars that are available in the local region of the dealer. In addition, we expect that local inventory has a direct price effect; consumers should be able to negotiate lower prices if all dealers in a region have high inventory.

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.

Third, we would like to calculate a first estimate of how much of the price variation in car prices can be explained with price discrimination through bargaining and how much can be explained with inventory fluctuations. In particular, we have empirical predictions about the relative magnitude of the inventory and bargaining effect 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 GM’s Saturn division and 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 Saturn and Autonation dealerships than other dealerships.

2.3 Estimation issues

Theoretically, 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 consumers' 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 results in low inventory. We would then find a correlation between low prices and low inventory. We weaken any such effect by measuring 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 previous 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 causes high prices). Of more concern would be endogeneity due to a demand shock. Suppose, for example, that there is suddenly increased consumer taste for a particular car. This would increase prices and run down inventories at the same time. Since we use car fixed effects, in order for this to be an empirical problem, the change in taste would have to be temporary. While examples of such short-lived changes of consumer tastes certainly exist, we don't believe they are common enough in the automobile

industry to substantially affect our results.

In general, any last-minute flexibility in resupply terms will reduce the estimated impact of days to resupply on price. We are in the process of gathering information about differences in nameplate re-supply policies. Our evidence so far suggests that dealers have little if any control over the timing of deliveries, and, close to the time of delivery, no control over the quantity delivered. Should there nonetheless be some variation in a dealer's ability to affect resupply in the short run, we expect that estimated inventory effects will be smaller for cars that are manufactured in the US than for cars that are manufactured abroad because the substantial shipping times from Europe, Japan, and South America decrease a dealer's ability to change the characteristics of his next shipment. We later estimate such a specification and find no evidence that estimated inventory effects are systematically smaller for domestic than foreign produced cars. Nonetheless, in a future version of this paper we plan to estimate an instrumental variables specification, with an instrument that is correlated with inventory levels but is uncorrelated with price. One possibility is precipitation. We expect variability in weather to cause short term shocks to inventory levels. Meanwhile, the assumption we impose on the data in our main specification is that the coefficient on the effect of inventory is constant across – possibly different – resupply schemes.

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, etc. Tracking inventory on the level of this definition is important because consumers have preferences over these attributes and some varieties of a make and model may be in short supply while the others are not.

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. This measure will be accurate from the beginning of our sample period: every car that was sold before the start of the sample is irrelevant for the dealer’s decision to price cars in inventory. Also, any car that was on the lot when our sample period began is likely to have been sold over the course of the 4 years of our sample, 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 end 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 all 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.

We try three different ways to scale inventory measures. First, we normalize inventory by average dealer sales volume. Second, we use indicators for when a dealership’s inventory is below certain percentile levels. Third, we limit the sample to dealerships which sell a minimum number of cars. The first approach is problematic in our sample because the variance of the denominator creates an unrealistically large variance in normalized inventory. The second approach is problematic because, given the fine granularity of our car definition, the 5th, 10th and even 25th percentile of inventory is likely to be 0 for small dealerships. This points to a larger problem, namely that there is not much variation in inventory of a particular car for small dealerships. Hence, we choose the third approach, namely to restrict the sample to dealership-car combinations for which the dealership sells at least 5 cars per month according to our definition of a car. This level is the mean across cars and dealerships and leaves 381,559 observations. Summary statistics for the dataset are in Table 1 in the appendix.

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 by 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.⁵

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, 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 basic 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 arrives on the dealership’s lot after having been on another dealer’s lot and perhaps already been test driven for some time. Also, a requested trade will have been driven from one dealership to the other. 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 nameplate 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*.⁶ We also use *TradeRequested* and *TradeReceived* as indicator variables in

⁵In 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.

⁶A 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.

subsequent price regressions. Since a dealer bears additional transaction and transportation costs for requested trades, we expect him to pass those on to the consumer.

For each dealer in the sample, we identify the maximum amount of inventory that dealer held over the three year sample. We create an indicator that is one when the total inventory held by the dealer is 95% or more of this maximum inventory. We name this variable *LotFull*.

To calculate the inventory levels across all dealers in the local area, we sum within a car, across all dealers in a DMA.⁷ This measure excludes the focal dealer’s own inventory. Not surprisingly, local inventory and dealer inventory are correlated. Hence, for use in some specifications, we use the residual from a regression of local inventory on dealer inventory; this creates the portion of the local inventory level that is uncorrelated with the focal dealer’s inventory (*eLocalInv*). Finally, we define a dummy variable that is one if supply is tight; *LocalInv25* is one when the inventory in a DMA for a specific car is below the 25th percentile.

We also define “dud” and “hot” cars using the average time such cars stay on a dealer’s lot across all dealerships in our sample. This measure proxies for the relative popularity of the car. 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 less than 22, we define the car as “hot.” If *DaysToTurn* is greater than 60, the car is classified as a “dud.” Both “hot” and “dud” cars encompass 25% of the sample. We call the middle 50% “regular” cars.

We define the end of a model year as the last 45 days before the introduction of the succeeding model. We may see different inventory effects during this time period, so in some specifications we allow some of the coefficients to take different values in this 45 day window. Additionally, it is possible to buy a car from model year t after cars from model year $t + 1$ are introduced. We omit these transactions from the dataset as their resupply conditions are not normal.

3.2 Dependent variable

The price observed in the dataset is the price that the customer pays for the vehicle including factory installed accessories and options and the dealer-installed accessories contracted for at the time of sale that contribute to the resale value of the car.⁸ 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

⁷DMA’s are metropolitan areas that correspond to TV markets, for example San Francisco-Oakland-San Jose, or Los Angeles.

⁸Dealer-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.

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.

3.3 Controls

We include a car fixed effect for each interaction of make, model, body type, transmission, displacement, doors, cylinders, and trim level. 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 vehicle from the average vehicle cost of that car in the dataset. This percent deviation, called *VehicleCost* will be positive when the car 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. Our measure of price also takes into account any variation in holdback and transportation charges.

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 inflation. If there are volume targets or sales on weekends, near the end of the month, or seasonally, we will pick them up with these variables.

We control for the number of months between a car’s introduction and when it was sold. This proxies for how new a car design is and also for the dealer’s opportunity 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 control for the competitiveness of each dealer’s market. For each dealership we count the number of dealerships with the same nameplate that fall in a zip code that is within a 10 mile radius of the zip code of the focal dealership. We control for cases where one owner owns several franchises. Hence, our measure counts only the number of separately-controlled entities.

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 the region in which the car was sold, either Northern or Southern California.

4 Results

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:

$$\ln(\text{Price}_i) = X_i\alpha + D_i\beta + S_i\gamma + \epsilon_i \tag{5}$$

The X matrix is composed of transaction and car variables: car, 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 S which contains various supply-side explanatory variables such as measures of inventory, days to resupply, dealer trade indicators, and indicators for when the lot is at capacity.

4.1 Estimation of the price-inventory relationship

We begin with an empirical test of the predicted effect of inventory and days to resupply on price. As previously discussed, we exclude from the sample all sales of cars which have been superseded by the launch of a succeeding model, as well as sales made by small dealers.

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.

Column one of Table 2 reports a simple test of the model where inventory and days to resupply are included linearly. Inventory has the expected negative coefficient; its magnitude, -.015, indicates that a dealer’s inventory must increase by 67 cars to raise price by one percent. The standard deviation of inventory in the sample is 18, so a two standard-deviation

change in inventory moves price by about one half of one percent. Days to resupply has an insignificant effect on transaction prices in this specification. To see if this result is driven by the restrictiveness of the linear specification, the next column reports the results of a spline in both inventory and days to resupply. We find much the same result when we use a spline for inventory, however, the marginal effects are somewhat larger; a change from one car in inventory to 21 lowers price by one percent. Increasing levels of inventory lead monotonically to lower transaction prices. Contrary to theory, however, the coefficients on days to resupply are negative in this specification. We will explore this finding further in the next subsection.

Other coefficient estimates have the expected signs. We find that consumers pay about 0.4% more, on average for a vehicle which was requested from another dealership. As expected, a larger number of competing dealerships within a 10 mile radius lowers the price paid by consumers. Demographic variables also have the expected sign but are not reported for the sake of brevity. 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.⁹

4.2 Effects by car supply, regional inventory, end-of-model-year and lot constraints

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

First, we distinguish between cars which are in short supply (“hot” cars) or in ready supply (“dud” cars) across all dealers in California. We expect to observe that inventory effects are larger for hot cars than for dud cars. This is because a dealer’s own inventory level is more of a binding constraint as excess demand for a model grows (see our discussion early in section 2.2). We replicate the last specification using interactions of our supply variables with hot and dud indicators; the results are reported in the last column of Table 2. Note that the specification is set up so that the main effects of inventory and days to resupply apply to all observations, while in the case of hot or dud cars, an additional coefficient must be added to the base effect. The results show that the main effect of inventory continues to hold for regular cars. The estimated coefficient is -.012, close to its previous level. We see strong effects according to whether a car is in net excess supply or not. The coefficient on inventory for hot cars has a much larger total effect on price, -.030% per car. This implies that a two standard deviation change in inventory level, *ceteris paribus*, would change the transaction price by one percent.

⁹For a thorough analysis of the effects of demographics on car prices please see Scott Morton, Zettelmeyer, and Silva-Risso (2003)

The effect of inventory on the price of dud cars is nearly identical to the effect on regular cars.

Dividing the effect of days to resupply into effects for hot, regular, and dud cars appears to partially explain its zero coefficient in the first specification. The coefficient on days to resupply for regular cars becomes .007 and significant in column three. A two-standard deviation change in *DaysToResupply* would alter price by 0.15%. Though this effect is positive as predicted by theory, it is empirically much less important than inventory levels. Again, resupply times have a bigger impact on the price of hot cars; the total estimated coefficient is for this group is 0.014, or double that of regular cars. The results show that the negative estimates of the *DaysToResupply* coefficients in the spline come from the pricing of dud cars. In this specification, we find that the total effect of days to resupply for dud cars is -0.011. The result implies that dealers reduce prices for cars that are in plentiful supply as resupply dates become more distant, all else equal. This general result persists throughout the paper and is robust to other functional forms we have tried. In contrast, we expected that days to resupply should not affect prices at all for cars in excess supply. We have no explanation for the negative effect of days to resupply on price for dud cars – it remains a topic of future research.

We have predicted that inventory should have a particularly large effect on price for transactions that take place near the end of the model year. The dealer’s opportunity cost of selling these cars decreases as the arrival of the new model draws nearer due to a drop in consumers’ willingness to pay for the old model. We repeat the last specification with the addition of interactions of our inventory and days to resupply measures with whether a car was sold in the final 45 days of the model year. The results are reported in the first column of Table 3. Both the effects of inventory levels and days to resupply change in the final 45 days of the model year. The main inventory effect increases from -0.011 to -0.018. The net coefficient on hot cars remains approximately constant at -0.028, while the net coefficient on dud cars increases dramatically to -0.031. Car pricing is more sensitive to inventory at the end of the model year, as predicted, and the increase is particularly sharp for cars in excess supply. The reason for this may be that dud cars are disproportionately harder to sell after the new model has arrived. The changes in the total estimated effects of days to resupply are less dramatic. The main *DaysToResupply* effect doubles during the end of the model year, but the effects for hot and dud are nearly identical to those that apply during the rest of the year.

We expect to find a stronger price effect of inventory when a dealer lot is close to maximum capacity. We add to the last specification the indicator *LotFull* that is one when the total inventory held by the dealer is 95% or more of his maximum inventory over the sample period. The estimated coefficient on *LotFull* is -1.34% (see column 2 of the table). This indicates that the average car sells for one percent less than usual when a dealer’s lot is full. The coefficients on our other inventory measures do not change appreciably.

Next, we add to the specification a measure of the inventory available at dealers in the same region, *eLocalInv*. As expected, we find that consumers who purchase a car when regional inventory is low pay higher prices (see column 3). The magnitude of the estimated coefficient indicates that price would increase 0.48% for a two standard deviation move in inventory *at other* dealerships. Note that the inventory coefficients at the focal dealership are similar to previous estimates, so the local inventory effect is in addition to the dealership specific effect. To more easily interpret the magnitude of the local inventory effect we repeat the specification with the indicator variable *LocalInv25* instead of *eLocalInv* in column 4. We find that that a consumer pays 0.4% more when the inventory in a DMA for a specific car is below the 25th percentile.

We conclude that these results are consistent with the predictions from the model: higher inventory results in lower prices (controlling for days to resupply), and a shorter time until the next shipment arrives (controlling for inventory) also lowers prices except for dud cars. The inventory effects are stronger for hot car models throughout the model year and for dud models close to the end of the model year. A full lot causes a dealer to lower price by about one percent, while a shortage of inventory of that particular car in the local area raises price by about 0.4%.

We end this subsection with some simple estimations of whether inventory effects differ by the country in which a car was manufactured. We expect that a dealer that can adjust its deliveries of cars in the short run will have a lower opportunity cost of selling a low-inventory vehicle car, and will thus show smaller inventory effects on prices; inventory would then be endogenous. We take advantage of a natural difference in manufacturer responsiveness to recent shocks to demand, namely distance between the manufacturing plant and US consumers. The endogeneity of inventory described above – should it exist – will likely be stronger for cars that are manufactured in the US than for cars that are manufactured abroad simply because shipping times cannot be altered. We run a specification where the interactions between our inventory measures and dud and hot are divided into cars manufactured domestically (including Canada) and cars manufactured abroad. We also include *LotFull* and *LocalInv25* in the regression.

We find no pattern of inventory effects being smaller for domestically-produced vehicles. The impact of inventory level on price is negative for domestic cars (-.016) and zero for foreign cars (see column one of Table 4). The coefficient on inventory for dud cars is slightly larger in magnitude for foreign cars compared to domestic cars, however the coefficients on inventory for hot cars are nearly identical. The coefficients on the *DaysToResupply* variables are all estimated at zero for foreign cars but follow the same pattern as previous estimates for domestic cars. Overall, we see *more* responsiveness of car prices to inventory for domestically-produced cars than for foreign-produced cars.

One explanation for this finding is that US car makers have different resupply policies than foreign manufacturers. For example, foreign manufacturers might be better at predicting consumer demand so that there is less variation in foreign car inventory levels. Hence, we repeat the specification restricting the sample to Asian car manufacturers, a group with more homogeneous (although not identical) resupply policies. There are both domestic and foreign produced cars in this subsample because several Asian manufacturers have production facilities in the US. The pattern of results is the same. Both inventory and days to resupply have stronger effects on the pricing of domestic regular cars than foreign regular cars (see column 2 of Table 4). Similarly, days to resupply for hot cars has a larger coefficient for domestic vehicles. Aside from these three differences, the coefficients are almost identical across the two groups. Interestingly, all the estimated coefficients on “dud” interactions in this specification are close to zero. This may indicate that some policy of non-Asian producers is driving the negative coefficients on *DaysToResupply* observed in previous results. Overall, this approach does not suggest that there is variation in dealers’ abilities to affect resupply in the short run.

4.3 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 inventory fluctuations relative to price discrimination. This is important because negotiated prices in car retailing are usually attributed to the fact that dealers try to price discriminate among consumers. We want to see how much of the price variation resulting from negotiated prices is instead driven by the dealer’s need to respond to inventory changes.

We know from prior work that prices vary substantially with demographics, even after controlling for region and dealer fixed effects (Scott Morton, Zettelmeyer, and Silva-Risso 2003). This indicates that our demographics are (imperfect) measures of the price discrimination effect due to bargaining.

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 demographic factors, the other the component of predicted price attributable to inventory factors. We calculate these two indices as follows: We run the last specification of Table 3 on a sample excluding no-haggle dealerships (see below). From the vector of estimated coefficients, $\hat{\beta}$, we extract two subsectors, $\hat{\beta}_d$ 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 subsectors and their corresponding data submatrices. Using the notation of equation (5), the demographic index is $\hat{I}_d = D\hat{\beta}$ where D includes income, race, home ownership, and all other consumer demographics. The inventory, or supply-side index is $\hat{I}_i = S\hat{\gamma}$, where S

includes inventory, days to resupply, end of model year, *LotFull*, local inventory, and whether the car was a dealer trade. 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 the inventory measures explain, on average, 21% of the combined inventory and demographic components of the predicted price.

Notice, however, that this number is a lower bound on the effect of inventory versus *price discrimination*. Our demographic measures pick up both price discrimination and variation in search techniques, outside options, etc., that are correlated with demographics. As noted in the introduction, with these methods we cannot separately measure how much price variation is due to price discrimination by dealers.

We end this section with an 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 and Saturn dealerships. AutoNation owns close to 300 dealerships nationwide and, like Saturn, has also adopted a policy of not negotiating with consumers. While there is some question as to whether individual dealerships always adhere to these policies, we expect inventory fluctuations in these dealerships to play a bigger relative role than in traditional dealerships.

We run separate regressions for transactions at AutoNation. 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 21% in the non-AutoNation dealerships compared to 51% in AutoNation dealerships. We repeat the same procedure for Saturn dealerships. Because there are only 4852 Saturn transactions in the dataset, we do not restrict the sample to dealerships which sell a minimum number of each “car” per month. As in the case of AutoNation dealerships, inventory fluctuations explain a larger fraction of the inventory and demographics based predicted price for Saturn dealerships (65%).

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 and Saturn 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 that demographics and price are correlated, we are able to show that inventory plays a larger role relative to demographics in the no-haggle setting.

5 Conclusion

We conclude from the evidence presented in the paper that local dealer inventory has a statistically and economically significant effect on the price at which new cars are sold in the United States. A dealership moving from a situation of shortage to one of ample inventory lowers transaction prices by about 1% *ceteris paribus*, or \$230 on the average car. Additionally, shorter resupply times also decrease transaction prices for cars that are not in excess supply. The total magnitude of the effect of supply side variables depends also on whether dealers face a lot size constraint, whether cars are at the end of their model year, and on the level of inventory in the surrounding DMA.

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.¹⁰ 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 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 21% in traditional dealerships. For so-called "no-haggle" dealerships, inventory explains 51% (for Autonation dealers) and 65% (for Saturn dealer) 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.

¹⁰carsdirect.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.

We conclude that price differences in car retailing are in part 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.

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.
- GALLEGO, G., AND G. V. RYZIN (1994): “Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons,” *Management Science*, 40(8), 999–1020.
- 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 (2002): “Coordinated Pricing and Production/Procurement Decisions: A Review,” mimeo, UC Berkeley, Berkeley, CA.

Table 1: Summary Statistics[†]

	Obs	Mean	Std. Dev.	Min	Max
Price	381559	24652.34	10137.78	5745.00	124945.00
Inventory	381559	17.64	17.58	1.00	80.00
DaysToTurn	381559	7.66	10.57	1.00	147.00
LocalInv	381559	81.37	102.80	0.00	930.00
LotFull	381559	0.02	0.14	0.00	1.00
TradeRequested	381559	0.07	0.26	0.00	1.00
Competition	381559	4.46	3.40	0.00	23.00
VehicleCost	381559	0.00	0.05	-0.50	0.61
SouthernCal	381559	0.55	0.50	0.00	1.00
%Black	381559	0.04	0.10	0.00	1.00
%Hispanic	381559	0.15	0.12	0.00	0.55
HispanicName	381559	0.13	0.34	0.00	1.00
%Asian	381559	0.11	0.12	0.00	1.00
AsianName	381559	0.08	0.27	0.00	1.00
Female	381559	0.32	0.47	0.00	1.00
Income	381559	6.32	2.64	1.13	15.00
Income ²	381559	46.98	39.76	1.28	225.00
%College	381559	0.34	0.17	0.00	1.00
%LessHighSchool	381559	0.12	0.12	0.00	1.00
%HouseOwner	381559	0.68	0.24	0.00	1.00
%Professional	381559	0.17	0.09	0.00	1.00
%Executive	381559	0.19	0.08	0.00	1.00
%BlueCollar	381559	0.25	0.15	0.00	1.00
%Technicians	381559	0.03	0.02	0.00	1.00
MedianHouseValue	381559	2.50	1.14	0.07	5.00
Age	381559	42.21	13.28	16.00	100.00
Age>64	381559	0.07	0.25	0.00	1.00

[†] MedianHouseValue in \$100,000. Income in \$10,000.

Table 2: Price effects of inventory[†]

Dep. Var. ln(price)	(1)	(2)	(3)
Inventory	-0.015 (0.001)**		-0.012 (0.001)**
DaysToResupply	0.001 (0.001)		0.007 (0.002)**
Inventory=1		0.975 (0.062)**	
Inventory=2-3		1.049 (0.039)**	
Inventory=4-6		0.543 (0.034)**	
Inventory=7-10		0.346 (0.031)**	
Inventory=11-20		0.198 (0.026)**	
Inventory>20		0.000 (0.000)	
DaysToResupply=0-2		-0.331 (0.049)**	
DaysToResupply=3-6		-0.206 (0.043)**	
DaysToResupply=7-12		-0.084 (0.041)*	
DaysToResupply=13-30		-0.369 (0.055)**	
DaysToResupply>30		-0.046 (0.083)	
Inventory*HotCar			-0.019 (0.001)**
Inventory*DudCar			0.001 (0.002)
HotCar*DaysToResupply			0.008 (0.003)**
DudCar*DaysToResupply			-0.018 (0.002)**
TradedCar	0.406 (0.030)**	0.348 (0.031)**	0.370 (0.031)**
Competition	0.005 (0.004)	0.001 (0.004)	0.006 (0.004)
Constant	998.616 (0.194)**	998.380 (0.199)**	998.596 (0.194)**
Observations	381559	381559	381559
R-squared	0.99	0.99	0.99

* significant at 5%; ** significant at 1%. Robust standard errors in parentheses.

[†] Unreported are all transaction characteristics, car fixed effects and demographics.

All coefficients are multiplied by 100.

Table 3: Price effects of end-of-model-year, full lot, local inventory[†]

Dep. Var. ln(price)	(1)	(2)	(3)	(4)
Inventory	-0.011 (0.001)**	-0.009 (0.001)**	-0.014 (0.001)**	-0.009 (0.001)**
Inventory*HotCar	-0.020 (0.001)**	-0.021 (0.001)**	-0.022 (0.001)**	-0.021 (0.001)**
Inventory*DudCar	0.004 (0.002)*	0.003 (0.002)	0.004 (0.002)*	0.004 (0.002)
DaysToResupply	0.004 (0.002)*	0.004 (0.002)*	0.005 (0.002)**	0.004 (0.002)*
HotCar*DaysToResupply	0.011 (0.003)**	0.012 (0.003)**	0.010 (0.003)**	0.011 (0.003)**
DudCar*DaysToResupply	-0.013 (0.003)**	-0.013 (0.003)**	-0.013 (0.003)**	-0.013 (0.003)**
Inventory*EndModelYr	-0.007 (0.002)**	-0.009 (0.002)**	-0.007 (0.002)**	-0.009 (0.002)**
HotCar*Inv*EndMdlYr	0.010 (0.004)*	0.011 (0.004)**	0.010 (0.004)*	0.012 (0.004)**
DudCar*Inv*EndMdlYr	-0.017 (0.004)**	-0.016 (0.004)**	-0.018 (0.004)**	-0.016 (0.004)**
DTR*EndMdlYr	0.019 (0.004)**	0.020 (0.004)**	0.019 (0.004)**	0.020 (0.004)**
HotCar*DTR*EndMdlYr	-0.020 (0.007)**	-0.020 (0.007)**	-0.019 (0.007)**	-0.020 (0.007)**
DudCar*DTR*EndMdlYr	-0.022 (0.006)**	-0.022 (0.006)**	-0.022 (0.006)**	-0.022 (0.006)**
LotFull		-1.336 (0.071)**	-1.316 (0.071)**	-1.330 (0.070)**
eLocalInv			-0.003 (0.000)**	
LocalInv25				0.392 (0.028)**
TradedCar	0.370 (0.031)**	0.372 (0.031)**	0.392 (0.031)**	0.385 (0.031)**
Competition	0.007 (0.004)	0.005 (0.004)	0.010 (0.004)**	0.008 (0.004)*
Constant	998.750 (0.199)**	998.724 (0.199)**	998.755 (0.199)**	998.619 (0.199)**
Observations	381559	381559	381559	381559
R-squared	0.99	0.99	0.99	0.99

* significant at 5%; ** significant at 1%. Robust standard errors in parentheses.

[†] Unreported are all transaction characteristics, car fixed effects, and demographics.

All coefficients are multiplied by 100.

Table 4: Price effects of inventory by origin of vehicle[†]

Dep. Var. ln(price)	(1)	(2)
DomesticInv	-0.016 (0.001)**	-0.020 (0.001)**
Hot*DomInv	-0.014 (0.002)**	-0.023 (0.002)**
Dud*DomInv	0.008 (0.002)**	0.019 (0.004)**
ForeignInv	0.001 (0.001)	-0.009 (0.002)**
Hot*ForInv	-0.031 (0.002)**	-0.038 (0.002)**
Dud*ForInv	-0.014 (0.003)**	0.009 (0.004)**
DomesticDTR	0.013 (0.002)**	0.009 (0.003)**
Hot*DomDTR	0.011 (0.004)**	0.017 (0.005)**
Dud*DomDTR	-0.030 (0.003)**	-0.013 (0.006)*
ForeignDTR	-0.001 (0.002)	-0.004 (0.003)
Hot*ForDTR	0.003 (0.004)	0.010 (0.005)
Dud*ForDTR	0.000 (0.004)	0.002 (0.004)
LocalInv25	0.375 (0.028)**	0.604 (0.036)**
LotFull	-1.238 (0.071)**	-1.816 (0.091)**
TradedCar	0.397 (0.031)**	0.733 (0.038)**
Competition	0.006 (0.004)	0.032 (0.005)**
Constant	998.358 (0.194)**	985.953 (0.257)**
Observations	381559	223629
R-squared	0.99	0.98

* significant at 5%; ** significant at 1%. Robust standard errors in parentheses. Second column restricts the sample to Asian nameplates only.

[†] Unreported are all transaction characteristics, car fixed effects, and demographics. All coefficients are multiplied by 100.