Wholesale Price Discrimination and Regulation: Implications for Retail Gasoline Prices

by

Justine Hastings

Yale University and NBER

<u>Preliminary Draft:</u>
<u>Please do not cite without permission.</u>

February 2009

I would like to thank Fabian Duarte for outstanding research assistance. I also thank Steven Berry, Philip Haile, Preston McAfee, Stephen Ryan, Sofia Villas-Boas, and Miguel Villas-Boas for helpful comments. I would also like to thank each of the individual gasoline dealers who made this study possible.

1. Introduction

The effects of wholesale price policies on retail prices and competition have long been a central topic of interest in industrial organization. The competitive effects of practices such as resale price maintenance, exclusive dealing, and wholesale price discrimination, are often ambiguous, with theories supporting both pro-competitive and anti-competitive reasons for instituting such practices. However, legislation regulating vertical contracts and restraints is often proposed as a mechanism to increase competition and lower retail prices. For example, the Robinson Patman act of 1936, makes wholesale price discrimination illegal if its effect is to harm competition.[1] It is rarely enforced since theory suggests that wholesale price discrimination, like retail price discrimination, may increase instead of decrease welfare. However, unlike retail price discrimination wholesale price discrimination could be used as a type of vertical restraint, lessening competition and increasing profits to the detriment of welfare.[2]

The empirical literature on wholesale price discrimination is relatively small compared to the literature examining retail price discrimination. This may be because wholesale transactions prices are not generally available or centrally collected, whereas retail prices are often visible, centrally stored and easier to obtain (for example from retail scanner data or hand-collected data from posted prices (Slade (1992), Noel (2007a,b), Verlinda (2007)). Thus, empirical research modeling wholesale price discrimination has depended on functional form and theoretical assumptions in order to infer consumer, retailer and wholesaler behavior and conduct policy simulations using only data on retail prices

[1] For a transaction to be in violation of the Robinson-Patman Act, sales of a commodity of like grade and quantity must be made at different prices to two or more buyers from a single seller with the effect of injuring competition (Act of June 19, 1936, 15 U.S.C. §§ 13-13b, 21a (2000)).
[2] For example, raising rivals' costs is a form of price discrimination. In addition, price discrimination can be used as a price floor in markets where the upstream firm sells both at retail and through a dealer. For a discussion of foreclosure and cost raising strategies and incentives for vertical integration, see Salop and Scheffman (1987), Ordover, Saloner and Salop (1990), Katz (1991), Hastings and Gilbert (2005) and Hortaçsu and Syverson (2007).

market shares and measures of input costs (Asker (2004), Villas-Boas and Hellterstein (2006), Villas-Boas (2007), Hellerstein (2008), Mortimer (2008), Villas-Boas (2009)).[3]

We collected a new data set on wholesale transactions prices, retail prices and volumes to understand the role of wholesale price discrimination in gasoline markets and examine possible implications of policies that eliminate wholesale price discrimination or regulate vertical restraints. While theoretical models imply that the competitive effects vertical restraints may be difficult to generalize across different markets (e.g. strategic complements versus strategic substitutes; Katz (1991), Rey and Stiglitz (1995)), gasoline markets themselves been the subject of proposed or actual regulation of vertical contracts and wholesale price policies.[4] In particular, legislative proposals aimed at banning wholesale price discrimination have been proposed in several states, and the Supreme Court has heard and ruled on cases involving violations of the Robinson Patman Act in gasoline markets.[5]

To examine the potential impact of eliminating wholesale price discrimination, we estimate a model of retail gasoline demand, retail and wholesale firm behavior using panel data on station-specific retail prices, retail volumes and wholesale transactions prices that we collected from gasoline stations in the San Diego, California metropolitan area.[6] First, we use the station-specific wholesale data to analyze the factors that drive price discrimination by branded refiners across stations and what determines retail mark-ups over wholesale prices. We find that refiners charge higher prices at stations which, based on observable characteristics, may face less elastic demand. We find that

[3] There are also several theory papers on wholesale price discrimination and vertical restraints, including Katz (1987), McAfee and Schwartz (1994), Yoshida(2000) and Villas-Boas (2009).

[4] Legislation that bans wholesale price discrimination has been proposed in all West Coast states as well as in New York, Connecticut, and Maryland. The legislation typically has different names. Common names include "Fair Wholesale Pricing Legislation", "Branded Open Supply Legislation", "Uniform Wholesale Price Legislation", and "Zone Price Elimination." For example, see the California State Attorney General's "Gasoline Task Force Report," and the Testimony Of Connecticut Attorney General Richard Blumenthal Before The House Judiciary Committee on April 7, 2000. http://www.naag.org/legislation/march/blumen041200.htm.

[5] For example, 110 S. Ct. 2535, Texaco Inc. v. Hasbrouck et al. (1990).

[6] We chose San Diego because of its tractable size, geographically isolated market location, and more importantly because we were able to enlist assistance from local trade organizations in our data collection efforts.

differences in wholesale prices explain nearly all of the observable differences in retail prices across stations of the same brand, and that retailers generally follow a constant-mark-up policy, setting retail prices as a constant mark-up over their station-specific wholesale price rather than marking-up prices in proportion to the inverse demand elasticity.

Next, we estimate a model of retail demand, wholesale and retail pricing that assumes upstream firms and vertically integrated firms play a static Nash-Bertrand game in differentiated products given the constant mark-up policy by retail dealers. Having the wholesale data to estimate dealer behavior allows us to test between alternative specifications of consumer behavior, instead of having to rely on demand assumptions to generate differences in price-cost mark-ups from which we infer retailer behavior (Nevo and Rossi (2008), Villas-boas (2009)). We estimate the model under two alternative parameterizations of the consumer's utility model; a vertical differentiation model of brand quality and a brand-loyalty model.[7] Both parameterizations generate reasonable station-level demand elasticities, with higher overall elastcities in periods of higher prices. However, when we compare predicted station prices generated by each parameterization to actual prices, we find that the brand-loyalty specification performs better than the vertical differentiation one. In both parameterizations, the estimated marginal costs of gasoline vary strongly the spot price of gasoline, and are generally lower for stations with alternative complementary retail products such as convenience stores.

We then use the estimates for the brand-loyalty model to simulate optimal prices if refiners were forced to charge one uniform wholesale price. We find that the market-wide-average average price of gasoline would rise by approximately five cents per gallon

[7] The former parameterization is prevalent in the applied industrial organization literature (Berry (1994), Berry, Levinsohn and Pakes (1995; 2004), Nevo (2000; 2001), Villas-Boas (2007)), with notable exceptions being (Berry, Carnal and Spiller (2006), Kalouptsidi (2008)), while the brand-loyalty model is often used in the marketing literature ((Chintagunta (1992), Chintagunta et al. (1991; 2001), Jain et al. (1994), Briesch et al. (2002))).

under our uniform wholesale pricing simulation.[8] We also find that total market-wide quantity sold decreases by 5.16%, indicating that welfare is lower under uniform wholesale pricing.

We demonstrate that branded refiner-marketers opt to raise the wholesale price of gasoline to their franchise dealers substantially under uniform wholesale pricing, while lowering the price through their vertically integrated stations. Thus they respond to regulation by shifting volume away from regulated stations towards ones where they have more discretion in price setting. These equilibrium effects imply that when we examine where our simulations predict prices will rise and fall as a result of uniform wholesale pricing, we do not always find that prices rise where reduced-form measures of wholesale prices were lower, for example. We often we find large price changes in places we might not expect to based on the reduce-form analysis, due to equilibrium interactions between strategic firms that can push product through dealer-run or vertically integrated stations. For example, stations in high-poverty neighborhoods do not have significantly different wholesale prices from other stations when controlling linearly for observable characteristics. However, those stations experience large price increases under uniform wholesale pricing, suggesting that uniform pricing may be regressive, raising prices more in low-income neighborhoods.

This paper proceeds in five sections. The first section provides a overview of the industry and the data used for analysis. The next two sections present the model and empirical results, and are followed by a section with policy simulations. The final section concludes.

2. Industry Background and Data Description

[8] It is important to note that our model does not include bulk sales to fleets or institutional buyers that may also impact the wholesale price if included in the uniform wholesale price rule. We also do not model the supply decision of branded versus unbranded wholesale gasoline as in Hastings and Gilbert (2005) or Hendricks and McAfee (2006), but instead assume that unbranded price of gasoline is a perfectly competitive price.

2.1 Background on gasoline prices and production

Wholesale gasoline is a commodity product and is supplied by firms that participate at all or only at some points of the production process, from crude oil extraction to refining to distribution and retailing.[9] In general, refiners produce wholesale gasoline and distribute it directly to retail gasoline stations or to intermediate wholesalers. Retail stations can have different contractual relationships with the upstream refiner-marketer, and wholesale price policies and distribution depend on that contractual relationship. Retail stations can be owned by a refiner-marketer or by an individual 'dealer.' If a retail station is owned by a refiner it can be operated by the refiner directly (i.e the refiner sets the retail price) or it can be leased to a residual claimant, a 'lessee-dealer'. In this case, the refiner will directly deliver gasoline to the dealer at a price determined on a daily basis for that station. This wholesale price, called the Dealer Tankwagon Price (DTW), is the subject of regulatory concerns over wholesale price discrimination since it can vary station-by-station even for stations that are in relatively close proximity. If a station is both owned and operated by a dealer, it can have a long-term affiliation with a branded refiner to carry and market their gasoline, and these 'contract-dealers' often must take direct delivery at a DTW just like lessee-dealer stations do.[10] If a dealer-owned station is unaffiliated with a refiner-marketer, it can purchase unbranded gasoline at a posted price from any supplier on the spot market or at a local distribution rack. Hence refiners cannot price discriminate between unbranded stations.

2.2 Data

To examine wholesale price discrimination and its effect on retail prices, we collected two unique data sets. Both data sets were collected for the market of San Diego, CA. We selected the San Diego because it has a relatively large share of stations paying DTW

[9] Gasoline of a particular grade meeting the same environmental regulations may be refined at different refineries and owned by different entities, but will be shipped comingled in pipelines and stored comingled in storage tanks for distribution to retailers and end users. In this sense, gasoline is a fungible commodity.

[10] In metropolitan areas, particularly those in the west coast, most stations are company operated or lessee-dealer formats, and almost all dealer-owned stations are contract dealers, paying a DTW. In rural areas and much of the gulf coast, for example, dealer-owned stations are supplied by intermediate wholesalers called 'jobbers' and pay a branded rack price set at the distribution each day by the refiner. The rack price is the same for all purchasers at that rack.

prices and a total number of stations that made collecting source data feasible. The first data set we collected was information on wholesale and retail prices from lessee and contract dealers throughout the greater metropolitan area ('dealers' from here on). We identified the set of possible gasoline stations and which stations where dealers using retail census data from MPSI Corporation, a company that collects information on gasoline station and convenience store locations and characteristics by metropolitan area for resale to industry participants. We augmented and cross-referenced these data with information pulled from refiner websites. Approximately 40% of the identified dealers in the San Diego area participated in our data collection efforts, supplying information on DTW's and retail prices on a weekly basis for eight weeks (April and May) in spring 2003 and six weeks (December and January) in winter 2003-2004.[11]

Second, we contemporaneously collected price and volume data for all stations (regardless of contract type) in the San Diego market. We were not successful in all circumstances, but were able to collect volume sold for about 85% of the 600 stations in the market. We sent a team of surveyors into the field to collect posted prices and record the volume sold at each pump on a rolling 3-week term. Since the volume of gasoline sold is recorded on the pump, recording the sales at two points in time and the prices posted at each of those two points in time gives a measure of total volume sold and average price during a three week period.[12] In the spring of 2003 we started by collecting prices and volumes for two of the three major areas in San Diego County (East County

[11] The local dealer trade organization ran an article describing our study and alerting dealers to the fact that we would be attempting to contact them to participate in our study. They reiterated that the study was for academic research, and asked dealers to consider participating in our study. This helped our recruitment efforts for the survey greatly. Each dealer was supplied with a confidentiality agreement stating that their individual identities would not be revealed, that their data would be kept confidential and used for academic research purposes only. They were allowed to choose if they would like a weekly contact by telephone, fax, or email to collect the pricing information. They received a small token of appreciation at the end of each survey round for their participation.

[12] We designed our own data collection since companies that collect price and volume data for industry participants typically do so in such a way that it is not useful for estimating demand. There are several such companies and they either survey station managers, asking them how much they typically sell in a month, or they collect volume information from gasoline pumps on a quarterly basis for a small but representative sample of stations in a market (for example a 20% sample). One exception is data from Kent Marketing used by Houde (2008), which contains actual volumes for all stations in Quebec City on a bimonthly basis (every 2 months).

and Central). Based on the success of the spring data collection effort, we expanded the coverage to North County as well for the winter.

Thus between the two data sets, we have prices and volumes for most of the stations, as well as DTW's for 40% of dealers in the San Diego Market. Figure 1 sets our sample time frame in the broader context of California gasoline prices from 2000 to 2006. Our spring sample covers a period of high but falling retail gasoline prices, while our winter sample covers a period of relatively low and stable prices. Thus price levels in our sample vary substantially over time, which will help us separate models of retailer behavior that differ in the prediction of how retailer margins should vary over time as price levels and price elasticities change.

We added to these data further information on key station characteristics from the MPSI retail gasoline census. The data from the MPSI station census include a geocoded location for each station, the traffic counts for each station on main and cross streets (which we will use to define markets and market size in the demand estimation), the size of the convenience store, if the station has a carwash, the number of fueling positions it has, whether if offers full service, and whether it has a service station.

Table I shows how the stations in our two samples differ from the population of stations in the market. Overall, we find that our sample stations are fairly representative of the population. The first column presents mean characteristics for the population of stations in the survey area (609), while the second column shows characteristics for the 547 stations from which we were able to collect price and volume data. The third column gives means for the sub-population of stations who are dealers, while the fourth column shows the same information for those dealers who provided us with DTW and retail price data.

In addition to the station brand and product characteristics, we constructed several local market characteristics. We used GIS software to construct measures of local competitive characteristics and local demographics by mapping each station and attaching to them their census block group, the set of census block groups with centroids within a 1.5 miles of them, and the driving distance between them and every other station in the market. We then merged on block group demographic data from the 2000 decennial Census for each block group with a centroid within 1.5 miles, and constructed population-weighted 'local' averages of block-group demographic characteristics. We also used the driving distances to construct local measures of competition such as the number of stations with one mile and the fueling-position-weighted HHI of stations within a mile. Again, the sample averages for these variables are consistent with the population averages.

3. General Model of Demand and Supply in Retail Gasoline

Our model of retail gasoline markets involves three types of players: consumers, gasoline retailers, and refiner-marketers, implying that we need a model of consumer choice, retailer pricing decisions, and refiner-marketer pricing decisions given consumer and retailer behavior. In this market upstream firms can both set wholesale prices for retail dealers who then set retail prices, and sell directly at retail through company operated stations (often competing with their own dealers at the retail level). There are also independent retail firms who procure wholesale gasoline in a competitive spot markets and compete in the retail market with branded refiners-marketers and branded dealers. We assume that firms set prices (retail or DTW) for stations in their portfolio to maximize profits as a function of demand, marginal cost, and dealer behavior (if they have dealers in the station portfolio) taking competitor's prices and characteristics as given. We use this model to generate moment conditions and estimate parameters of the utility and marginal cost functions using aggregated price and market share data following the methodology developed in Berry, Levinsohn and Pakes (1995).

3.1 Consumer Behavior

First, we assume that individual consumers choose to purchase from the gasoline station that maximizes utility as a function of prices, station characteristics, and preferences.

$$u_{ijt} = X'_{jt}\beta_i - \alpha_i r_{jt} + \xi_{jt} + \varepsilon_{ijt} \qquad (1)$$

where:
X_{jt} = a *Kx1* vector of station characteristics at station *j*
r_{jt} = the price for gasoline at station *j*
α_i = consumer *i*'s price sensitivity
β_i = a *Kx1* vector of consumer *i*'s taste or preference for station characteristics
ξ_{jt} = mean value of unobservable (to the econometrician) characteristics of station *j*
ε_{ijt} = mean zero stochastic term

Given this utility, the market share of each firm is given by the integral over all consumers in the market of the probability that consumer *i* chooses firm *j*.

$$s_{jt}(p, X, \xi; \theta) = \int_{A_{jt}} dP(\varepsilon) dP(\beta) \qquad (2)$$

where $A_{jt} = \{\beta_i, \varepsilon_{ijt} \mid u_{ijt} > u_{ikt} \forall k \neq j\}$.

Note that this specification abstracts from search costs (Sorensen (2000), Lewis (2005) and (2008), Chandra and Tappata (2008)), assuming instead that agents know the characteristics of products in their choice set. It also does not allow consumers to vary in the decision over how many gallons to purchase and abstracts from the choice of which grade of gasoline to purchase given the decision to purchase at the station. This is a reasonable and perhaps necessary assumption as the relative prices across grades of gasoline within a station are almost always constant over time, with the station typically setting the price of regular grade gasoline and marks-up mid-grade and premium by a standard 10 cents each.

3.2 Cost Function

We assume that there are *F* firms in the market, each of whom derive profits from some subset, F_f, of the *J* stations. For a multi-product refiner-marketer this subset of stations would include both their directly operated retail stations as well as their dealer-operated stations. For an independently owned unaffiliated station, there will only be one station in the firm's profit function.[13] We assume that the marginal cost of supplying gasoline at each station is a linear function of factors that affect the per-gallon cost of gasoline, such as the spot price of gasoline, seasonal effects, as well as station characteristics that may shift the marginal cost of gasoline such as convenience store size, whether or not the station has a service bay, full service fueling, and the distance from the wholesale distribution rack to the station.[14]

$$mc_{jt} = \eta W_{jt} + \omega_{jt} \qquad (3)$$

3.3 Dealer Behavior

Each station that is dealer-run has two parties that sequentially set prices to determine the retail price. The refiner-marketer sets a DTW for each station, and the dealer decides the retail price given demand, competitor prices, and the DTW. We define the function $r_j(\cdot)$ as the function that transforms the upstream firms price, p_j, into a retail price at station *j*. For a company-operated station or unbranded station, this function is $r_j(\cdot) = p_j$. The upstream firm is the downstream firm, so there is no second mark-up. For a dealer-run station, this function could be the solution to a Nash-Bertrand game if the dealer has local market power, or it could be $r_j(\cdot) = p_j + m_j$ if the station follows a simple rule-of-thumb market or behaves as if it is in a perfectly competitive market. We will determine the form of this function in next section, when we use our panel of DTW data to examine the pricing rule that dealers follow.

3.4 Pricing Problem for the Firm

Given demand, each firm acts to maximize profits in a static Nash-Bertrand game with differentiated products. There are *F* firms in the market. Each firm *f* can be an integrated

[13] It is possible for an individual proprietor to own multiple stations in one area. We do not have information on which independent dealers own which stations and lease which branded stations, so we assume each unbranded station is a one-station proprietor and that each lessee dealer only leases or owns one station.

[14] In reality, the gasoline station should be modeled as a multi-product firm that sets prices to jointly maximize profits from gasoline sales as well as sales from ancillary product services such as convenience store goods and car washes. Since I do not have data on convenience store products, prices, and volumes sold, a reduced-form way to incorporate the effects these products have on gasoline prices is to allow them to shift the marginal cost of gasoline.

refiner owning multiple lessee-dealer or vertically integrated stations, or it can be an independent retailer with only one station under management. The profit function for firm *f* is given by (from this point on we suppress time subscripts):

$$\Pi^f = \sum_{j \in F_f} \left(p_j - mc_j\right) s_j\left(r\left(p_j, r(p)_{-j}, X, \xi; \theta\right), r(p)_{-j}, X, \xi; \theta\right) \qquad (4)$$

where p is the vector of prices in the market and F_f is the set of stations that firm *f* derives profits from. Given the profit function, the firm *f*'s first-order condition for price at station *j* owned by firm *f* is given by:

$$\frac{\partial \Pi^f}{\partial p_j} = s_j\left(r(p)\right) + \sum_{k \in F_f} \left(p_k - mc_k\right) \frac{\partial s_k}{\partial r_j} \frac{\partial r_j}{\partial p_j} = 0 \qquad (5)$$

This profit condition must hold for each firm *f* for all of its stations *j*.[15] Hence the first-order conditions implicitly define optimal prices for each station given demand parameters, vertical structure of the station, retailer pricing behavior, and marginal costs.

Following Berry, Levinsohn and Pakes (1995), we use the exclusion restriction that the structural error terms are orthogonal to a set of instruments to form moment conditions to estimate the parameters of the model, using GMM.

$$E\left(\xi_j \middle| Z\right) = E\left(\omega_j \middle| Z\right) = 0 \qquad (6)$$

where ξ_j and ω_j are functions of the parameters of the model and the data that defined by the utility function and the firm's first-order conditions combined with the marginal cost equation.

[15] Note that we are not modeling the unbranded market as part of the profit function for integrated refiner-marketers as in Hastings and Gilbert (2005) and Hendricks and McAfee (200X). In effect we assume that the spot price of gasoline is set competitively so that integrated refiner-marketers cannot affect the input cost to independent downstream rivals.

Once we have parameter estimates for demand, retailer behavior and marginal costs, we can implicitly solve for optimal prices under price discrimination or uniform wholesale pricing. Equation (5) gives the first order condition in the market assuming that firms can set prices at each station independently. Under a uniform wholesale price regulation, any prices charged to dealer stations must be the same, implying that the first order conditions now change. The firm will solve the above first-order condition subject to the constraint that $p_j = \overline{p}^f, \forall j \in D^f$, where D^f denotes the set of firm *f's* stations that are dealers. Thus if a firm is a single-product firm or a refiner-marketer with all directly operated stations, the regulation would not affect the firm's first order condition directly. However, if a refiner-marketer had *J* stations, all of whom were dealers, then firm *J* would go from maximizing profits over *J* prices to maximizing profits over one single price. With demand, retailer behavior, and cost estimates, we can implicitly solve for optimal prices subject to the uniform wholesale pricing constraint.

4. Estimation and Identification

4.1 Estimating Retailer Behavior

We begin by examining the panel of station-specific wholesale prices to see how retailers respond to changes in DTW prices, and how DTW prices respond to changes in spot prices to see if the raw data appear to reject or support particular models of retail behavior. Potential models could include perfect competition, competition on price in a differentiated products market, or nonstandard models such as rule-of-thumb mark-ups.

Figure 2 plots weekly average retail prices, average DTW prices for the 40% sample of dealers in our survey along with the weekly average spot price of gasoline in Los Angeles.[16] All taxes have been removed from the retail prices in the graph to facilitate comparing retail prices and DTWs to spot prices. There is a 7.75% sales tax and \$0.348

[16] We purchased LA Harbor spot prices from Oil Price Information Service (OPIS).

in excise taxes. The plots are made separately for the spring and winter survey rounds. The graphs imply that regardless of the overall price level, retail prices appear to have a constant mark-up over DTW. This average mark-up is about 10 cents per gallon. In addition, the graphs imply that changes in DTW's are passed completely through to retail prices.[17]

Figure 3 plots average retail mark-up over DTW cost as well as DTW mark-up over wholesale spot price, along with 10th and 90th percentile bands around each mean. DTW mark-up over the spot price varies substantially over time, with DTW prices consistently 50 cents a gallon higher than spot price for the first part of the spring round. In addition, DTW prices can vary substantially across stations, with a spread of 10-15 cents per gallon. The DTW spread is not always constant, suggesting that DTW's may be optimized by station over time to adapt to local intertemporal market conditions. On the other hand, retail mark-ups over DTW, while they have a spread across stations that is as large as their mean, appear to stay stable at ten cents per gallon across the spring and winter rounds with a stable spread around that mean.

Table 2 presents OLS and IV regressions of retail prices net of taxes on DTWs. The first column shows that, on average, retail prices rise by 0.983 cents when DTW increases by $1.00. This point estimate is not significantly different from 1, and DTW and a constant explain 97% of the variation in retail prices net of taxes. The second column instruments for the DTW using the LA harbor spot price of gasoline and its interactions with seasonal dummies, the distance from the station to the distribution rack, and brand dummies. The coefficient on DTW gets slightly closer to 1 (0.992).

The mean and median station-level mark-up are 10.0 and 10.8 cents per gallon respectively. The average mark-up at a station ranges from about 6 cents per gallon to 20 cents per gallon. So while the mark-up level varies across stations, within station is it constant over time and does not vary with the overall price level in the market. In fact, if

[17] Borenstein, Cameron and Gilbert (1997) and Borenstein and Shepard (2002) examine asymmetric price adjustment between retail and rack prices and rack prices and crude oil prices in gasoline markets.

we regress the retail mark-up over DTW on the average retail price in the market, controlling for station-level fixed effects, we find an insignificant effect of gasoline price levels on retail markup, with a coefficient near zero (-0.008). This suggests that either retailers have fairly large differences in marginal retailing costs and price at marginal cost, or that they capture rents through a constant mark-up even if they don't adjust those rents as price levels and demand elasticities fluctuate. The latter would be consistent with a rule-of-thumb mark-up policy where retailers set their rule at an average optimal mark-up given market power and typical gasoline price levels and fluctuations.[18]

There is significant variation in DTW's across stations in addition to variation in the amount retailers mark-up over DTW. Table 3 examines which factors affect the relative DTW stations pay by regressing DTWs on observable station characteristics, local demographics and local competitive characteristics. All regressions include weekly fixed effects with standard errors clustered at the station level. The first column has DTW as the dependent variable, while the second column has retail price *less* DTW and taxes as the dependent variable. Hence, significant coefficients in the second column represent factors for which DTW variation does not completely explain variation in retail prices.

Looking first at station characteristics, each of the refiner-marketeers charges a higher DTW to their stations than ARCO, the excluded brand. However, retail mark-ups are not significantly different across brands implying that brand-average differences in retail prices are the results in differences in DTW prices across brands rather than differences in retailer margins, all else equal. Stations that are located on freeway exits and stations that are longer distances away from the distribution rack pay higher DTW's as well. Interestingly, retail stations on freeway exits also charge higher retail prices implying that refiners only capture 1/3 of the overall retail price premium at freeway-exit stations. A similar pattern arises for stations offering full-service; in this case variation in retail prices is generated by additional retailer margins instead of increased DTW mark-up.

[18] During the data collection process, many dealers actually told us that they followed a rule-of-thumb mark-up and this rule-of -thumb varied across dealers, and within dealer across stations that they ran in the handful of cases where a dealer ran more than one station.

This is intuitive in this case if offering full-service increases the actual marginal cost of the retail station operator.

Local demographic characteristics are significant determinants of DTW prices. Local demographics are calculated as population-weighted averages of 2000 decennial Census block group data for block groups with centroids within 1.5 miles of a station's location. Refiners charge higher average prices in higher income neighborhoods, but lower prices in neighborhoods with a larger fraction white, with a longer average commute time, and a larger number of cars per house (suburbs). In addition, they charge significantly lower prices in markets with a higher number of stations with the same brand. They also charge significantly lower prices in markets with unbranded stations, although the effect of an additional unbranded station is half of the size of the effect of an additional station with the same brand. Looking at the retail margin column, none of these variables are significant explanatory variables for retail margins, implying that wholesale price discrimination is the primary factor explaining retail price variation across these dimensions.

Figures 2 and 3 and Tables 2 and 3 suggest that retailers tend to pass through changes in DTW's completely, and charge a constant mark-up over DTW that varies with station characteristics. We will show in the next section, that demand elasticities imply that retail stations have local market power, and that demand elasticities are higher in times of high prices. The mark-up levels and the lack of correlation with price levels over time is inconsistent with a model of double-marginalization by retailers who set prices to maximize profits in a Nash-Bertrand game given costs and competitor's prices. It is consistent with a rule-of-thumb mark-up or competitive pricing where differences in mark-ups are generated by differences in approximate optimal prices or differences in marginal costs. To estimate dealer behavior, we therefore predict mark-ups for all dealers as a function of station-level characteristics estimated off of our sample of dealers who provided us with DTW and retail price data.[19] Thus, $r_{jt}(p) = \hat{\alpha}_j + \widehat{\beta}_j p_{jt}$. We will plug these dealer behavior functions into the firm's first-order condition for firms that have dealers in their station portfolio, so that firms set DTW prices and retail prices at their company-operated stations taking this predicted dealer behavior into account.[20]

4.2 Parameterizing the distribution of consumer preferences

There are two key types of parameterizations that come up in the literature on estimating demand for retail products. The most common one assumes that products are vertically differentiated, with customers drawing preferences for brands and other quality attributes from independently, identically distributed distributions (for example, see Berry (1994), Berry, Levinsohn and Pakes (1995; 2004), Nevo (2000; 2001), Villas-Boas (2007)). Alternatively, several researchers have used discrete type models where customers appear in latent or 'types' in the market place, and each type has a set of degenerate preferences, creating correlation between the idiosyncratic preferences for product attributes (Berry, Carnal and Spiller (2006), Kalouptsidi (2008)).[21] Preferences characterized by brand loyalty fall into this category, and this parameterization has been used to estimate demand in many retail goods markets where a handful of dominant brands compete with generic products (for example, Chintagunta (1992), Chintagunta et al. (1991; 2001), Jain et al. (1994), Briesch et al. (2002)). Often, these papers use individual-level scanner data and estimate brand-loyal type probability using past purchase behavior. We will only have station-aggregated data in this case, making it more difficult to separately identify type probabilities from their distinct preferences over brands. Several papers have found reduced-form evidence consistent with brand-loyalty in gasoline markets (Hastings (2004), Lewis (2005; 2008)), and this preference parameterization may produce different

[19] We use as explanatory variables the station's brand and its interactions with other station characteristics and local demographic and market structure characteristics to predict retail mark-ups, providing a reasonable prediction of mark-ups by station.
[20] We also estimated the full model under the assumption that dealers are Nash-Bertrand players and find as expected that the retail prices predicted by that model are much higher than actual retail prices, and reject that model under each of the preference parameterizations we consider. Dealer mark-ups are too low given demand elasticities to be the result of a second optimal mark-up.
[21] The nonparametric estimator proposed by Bajari et al. (2008) can allow for both of these parameterizations. The types model is a mixture of distributions; degenerate in this case but it could be a mixture of normals, for example, each occurring with a probability in the population.

substitution patterns and predictions across local markets where different brands compete with each other and with unbranded stations.

Since the focus of this paper, and of structural estimation more generally, is on prediction. So we will estimate a handful of models with these different parameterizations and use the one that has the best in- and out-of-sample predictions to conduct our counterfactual simulations.

The general random utility model for estimation is given by:

$$u_{ij} = X_j^{1'}\beta^1 + B_j^{'}\beta_i^b - \alpha_i p_j + \xi_j + \varepsilon_{ij} \qquad (7)$$

X_j^1 includes station characteristics which are assumed to have constant preferences in the population: whether the station is located on a freeway exit, if the station has a carwash, if the station has a service station, if the station has a large (greater than 600 sqft) convenience store, and the number of fueling positions the station has, and if the station is primarily a convenience store with little selling space for fuel. B_j is a matrix of brand dummies for each of the five major branded refiner-marketers: ARCO, Chevron, Mobil, Shell/Texaco, and Union 76. Preferences for prices and brands are allowed to vary idiosyncratically in the population, and the parameterizations of these idiosyncratic preferences distinguish the vertical differentiation from the brand loyalty models.

Specifically, we assume the following:

Vertical differentiation model:

$$\begin{bmatrix} \alpha_i \\ \beta_i^b \end{bmatrix} = \begin{bmatrix} \alpha \\ \beta^b \end{bmatrix} + \Gamma D_i + \Sigma \nu_i \qquad (8)$$

Where D_i represents consumer demographics; draws from the distribution of local income local commute times using block-group data from the 2000 decennial census, and υ_i are independent draws from a standard normal distribution with variance parameters of diagonal matrix Σ to be estimated.[22] We assume that ε_{ij} are distributed *i.i.d.* extreme value, yielding the familiar mixed-logit discrete choice model.

Brand Loyalty Model:

The brand-loyalty model makes the assumption that there is a latent brand-loyal customer for each of the major brands in the market. There is also a 'shopper' who has no brand loyalty and values all gasoline the same (zero brand preferences). Each brand loyal type has an idiosyncratic preference for its brand over all others, and these types occur latently in the population with probability, π_t, to be estimated. Thus for the brand-loyalty model, $\beta_i^b = \Gamma D_i$, where D_i is an Nx1 vector indicating which brand-loyal type the customer is and N is the number of brand-loyal types, and Γ is a diagonal NxN matrix where the diagonal elements are the values that each type places on its own brand. The D_i are drawn from a discrete distribution, so that each type occurs with probability π_t, such that $\sum_{t=1}^{N+1} \pi_t = 1$.

Thus, conditional on α_i, the market share for each station will be the probability-weighted average of conditional logit shares for each type of customer (Berry, Carnal and Spiller (2006)). We allow preferences for price to vary idiosyncratically with the distribution of income and commute time in the population just as in the vertical differentiation model, so that $\alpha_i = \alpha + \Lambda D_i + \Psi \nu_i$, as described in the vertical differentiation model above.

4.3 Defining markets and market shares

To construct market shares from our utility model, we first need to define what a market is. Gasoline, like other products such as movie theaters (Davis (2006)) or funeral homes (Chevalier, Harrington, and Scott-Morton (2008)) is a geographically differentiated

[22] To simulate income and commute time characteristics we took the set of census block groups with a centroid within 1.5 miles of any of the stations in a market. For income, we then fit a log-normal to the distribution of income levels in each market, and we take draws from those log-normal distributions to simulate individual income. For commute time, we drew from the empirical distribution of commute times in the local market.

product. However, unlike movie theaters and funeral homes, gasoline stations are not 'destination products', meaning that people tend to fill up their tanks on their way to something, and do not typically drive from their home to a station and back again for the purpose of buying gasoline. Thus while prior studies have typically simulated customers using distribution of residential locations, it is not clear that is the appropriate model here. One approach is to follow Houde (2008), who estimates a structural model for gasoline demand in Quebec City, by simulating customers on shortest routes from homes to offices using residential and work location information from a Canadian Census. An alternative is to define local markets and use traffic count data, assuming that each car that passes through a geographic region can decide to purchase gasoline at one of the stations, or continue driving. The former case has the benefit of using actual commute patterns. The latter case does not restrict people to be looking for gasoline only on work-home trips, and better exploits differences in retail characteristics and brand concentration across local markets to identify heterogeneous preferences for brands and prices.

In addition, the topography of San Diego metropolitan area lends itself nicely to defining local sub-markets. Figure 4 shows a graph of the metropolitan area with stations locations plotted. We used the geocoded station locations to calculate driving distances from each station to every other station. We then used an algorithm that looked for gaps in the distance from each station to the next nearest station to define markets as sets of stations that are close to each other, and have a significant driving distance gap to the next nearest station not included in the market. We tried a few different algorithms, and test which one produced price predictions that fit the data best under both vertically differentiated and brand-loyal models. We will present results based on the best-fit market definition which has around 60 markets. Conversations with the consulting company that constructs DTW price-discrimination zones for the major oil companies suggested that the largest companies define about 50 separate markets in San Diego.

Once we have defined markets, we use traffic count data from the MPSI station census to construct the size of the market and thus market shares. MPSI records the number of cars passing by a station on both primary and secondary streets, so we use this information to calculate the total number of cars passing by the stations in each market and thus the total market size. The outside good is purchasing at another station in another market, or not purchasing gasoline at all.

4.4 Instruments and identification

We estimate the parameters of the model using GMM, creating moment conditions from orthogonality conditions between the structural errors from the market share equation and the marginal cost equation and a set of instruments (Berry (1994) and Berry, Levinsohn and Pakes (1995)). Our instruments include traditional cost shifters, the spot price of gasoline and the distance to the gasoline rack. We also include competitor characteristics and their interactions with the spot price as instruments. In particular, we include interactions between a station's own brand and the brand identities of its competitors, and interaction these with the spot prices, since our two preference parameterizations have different predictions about how intensely branded stations compete with each other and with unbranded stations (Hastings (2004)). We also include other competitor characteristics such as number of fueling positions, whether they have a carwash, offer full service, or have a service bay.

5. Results

Table 4 presents parameter estimates for the first specification using the vertical differentiation model. It includes only demographic draws from the local income distribution. The results are generally intuitive. The mean coefficient on prices is negative and significant. Having a carwash significantly increases demand at a station, while stations with service bays have significantly lower demand. Having many fueling positions and being located off a highway exit both lead to significantly higher demand. A dummy for "7-Eleven" convenience stores is included because this convenience store chain typically does not offer gasoline, and when it does, the stations are typically old with stations with only 1 or 2 fueling positions, so their primary operation is as a convenience store. They sell lower volumes than most unbranded stations, and the chain ended up closing many stations after our sample period, so the model with an indicator for this station type fits better than one that excludes it. The coefficient is large, negative and significant, implying that customers value these stations significantly less than other unbranded or branded stations. Mean brand preferences for the major refiners are positive and significant with the exception of Mobil.

The idiosyncratic preference for price varies positively and significantly with income, implying that consumers from higher-income neighborhoods are less price sensitive. Income is measured in tens of thousands of dollars, implying that a 25,000 increase in income would move the preference for price up by 0.225, which is about a 8.7% change off the base of -2.578. Idiosyncratic preference for Arco is increasing in income, while the preferences for Shell and Unocal decrease significantly with income. For an average person with $46,000 in family income, Shell is the preferred brand, followed by Arco and Chevron while Mobile and Unocal are not significantly valued over unbranded stations. In general, the standard deviation estimates from the draws from the normal distribution meant to simulate unobserved tastes are close to zero and highly insignificant, suggesting either that preferences vary close to linearly with income, or that market share data alone are insufficient to precisely identify these additional parameters (Berry Levisohn and Pakes, (2004), Hastings Kane and Staiger (2007)).

Table 4 also presents the computed demand elasticities implied by the demand estimates along with their standard errors. Standard errors were computed analytically using the delta method. Average demand elasticity is highest at -3.535 during the highest-price time period 1, and falls to a low of -2.989 in the lowest price period 4. The spread from the 90th to the 10th percentile also decreases from (-4.324, -2.874) in period 1 to (-3.472, -2.475) in period 4. In all periods all stations have demand elasticities that are negative and significantly different from zero (as demonstrated by the 90th upper confidence interval for the demand elasticities).

Table 5 presents parameter estimates for the second specification of the vertical differentiation model which additionally allows preferences to vary with commuting time. The coefficients on variables without random coefficients are well identified and generally similar to those in the prior table. The mean preferences for brands, however are quite different with even larger standard errors than the prior specification. The interaction between price and income remains significant and quite similar to that in Table 4. The interactions with brands and income are also still negative, and remain generally significant. However, none of the interactions with commute time are significantly different than zero, including the interaction with price, and the variance in idiosyncratic preferences remain generally close to zero and highly insignificant. Adding commute time also yields large standard errors on the demand elasticities. Looking at the estimates in Table 5, we see again that demand elasticity is larger in times of high prices, and falls as price levels drop. And while these estimates have similar means, almost none of them are significantly different from zero as can be seen from the upper confidence intervals.

Table 6 presents results from the brand loyalty model. The mean preferences for price is negative and significant as in Tables 4 and 5. The mean coefficients for ancillary products like carwash, service bays, number of fueling positions and “7-11” are similar in sign, magnitude and significance to those in the vertical differentiation model. The random coefficients on the brands are now summarized by the share of each type in the population and the value they place on their brands. The mean values for the type-shares imply that Chevron types have the largest share, followed by Arco, Mobil and Shell. The share of Unocal types is not significantly different from zero. Chevron, Shell, and Unocal stations all charge relatively high prices in the market, but Unocal volumes are lower on average. In addition, Arco has a significant share of loyal types who have a significant positive valuation of the brand, and while they typically charge lower prices than Chevron and Shell, their volumes are larger on average. Similar to the vertical models, Unocal and Mobil are the lowest brand values based on their combination of their type's shares and the valuation their loyal customers place on their brand over all others.

The residual share of consumers in the market, 25.2%, is the "Shopper" type who values all gasoline brands the same. Thus the demand elasticity a station faces depends on its brand and the brand identities of the other stations, and a market with few branded stations may result in more elastic demand all else equal. Also in this specification indicates price sensitivity decreases significantly with both income and commute time, implying softer competition in high-income high-commute markets even if there are few strong-brand competitors present. An increase in income of $25,000 implies a decrease in price sensitivity of 7.1%, and a 10 minute increase in commute time would imply a 2.2% decrease in price sensitivity. Thus a market with only a couple of brands present but located in a high income may not be as price competitive. A person with average income of 46,000 and an average commute time of 25 minutes would have a coefficient on price of -2.166 (4.6*0.0751+25*0.006+-2.661). On average, demand is slightly more elastic in each time period in the brand-loyal specification than in either of the vertical specifications, and all of the demand elasticity estimates are significant at the one percent level or higher.[23]

Table 7 shows the estimated parameters from the marginal cost function given in (5). The parameter estimates are very similar across the three models. In all three models the marginal cost of gasoline increases 1 for 1 with the spot price of gasoline measured as the gulf coast spot price of federal reformulated gasoline. In addition, many of the ancillary products and services have effects on marginal cost with signs and magnitudes that seem reasonable. As stated before, since we lack data to estimate a full model of a multiproduct retailer, on approach to include profits from complementary products is to allow them to shift the marginal cost function for gasoline. A car wash could either increase or decrease the marginal cost of selling gasoline, and the vertical models show that it increases the cost. Convenience store products, since they typically are priced with large margins should decrease the marginal cost of gasoline as long as people are more likely to buy products in the convenience store when they buy gasoline. All three models have a negative coefficient for this variable. The coefficient on the distance to the distribution rack has the opposite sign than expected. Stations further from the rack should have higher delivery prices per gallon of

[23] Note that standard errors of the parameter estimates in all three of the models improved when we estimated the model using moment conditions implied by the demand and supply equations instead of just using those generated by the demand side. In particular the brand loyalty estimates were very imprecise without the additional moment restriction from the supply equation.

gasoline, however the estimates are decreasing in distance from the rack. This may be because gasoline is delivered along routes, and we are not capturing the true distribution costs since that depends on how many other stations of a particular brand are located near a station, not just the distance that station is from the rack. We also allow the marginal cost of gasoline to vary seasonally, including time dummies for each period in the data set. In theory, all variation in marginal input costs should be captured by the spot price of gasoline, however, as is apparent from Figure 1, the spot price of gasoline is very volatile in California and retail prices do not always follow the spot price. Margins between retail and spot price are also a function of gasoline and blending component stocks when estimated over a longer panel than available in this data set. Since these stock data are available only on a monthly basis, including them is equivalent to including time dummies. Without a longer time series it is difficult to fit the correct effect of stocks on marginal cost, so we include the time dummies instead.

6. Policy Simulations

6.1 Uniform Wholesale Pricing

We can now use our model estimates to distinguish which of the models is able to predict prices in the current regime most effectively, and then simulate prices in a counterfactual world with uniform wholesale pricing. To do so, we first use the first-order conditions to predict prices in the current price discrimination regime.

Figure 5 presents kernel density estimates of the prediction error, predicted – actual price, across the three models. The brand-loyalty model has the best fit, and the sum of squared errors is half the size for this model than for the other two.[24] Given the fact that the brand-loyalty model has a better prediction for the actual prices in the market, we will use this model to predict the impact of uniform wholesale price regulation on retail prices. In order to do so we solve the system of first-order conditions restricting refiner-marketers

[24] SSE's are 13.08, 16.10, and 36.94 for the three models respectively.

to charge the same price to all of their dealer stations in each time period across all markets.

The mean price under price discrimination of 1.793 rises to 1.839 under uniform wholesale pricing.[25] Thus banning price discrimination leads to an increase rather than a decrease in average prices. Figure 6 presents kernel density estimates of the distribution of predicted prices under price discrimination and under uniform wholesale prices. Note that a large fraction of stations that were mid-priced in the distribution now shift substantially to higher-priced territory, while there is a small growth in density for very low priced stations. Thus banning wholesale price discrimination in effect increased rather than decreased the variation in prices. The reason this occurs is because integrated refiner-marketers have both dealer run and vertically integrated stations. When they are forced to choose one price for all dealer-run stations, they choose a price near or at the top of the distribution of their DTW's under price discrimination, and lower prices at their vertically integrated stations.

Table 8 shows the mean change in predicted prices and volumes by brand and contract type when uniform wholesale pricing is in place, controlling for time period fixed-effects. The first column shows that all branded refiners raise DTW's substantially, causing prices to increase significantly at their dealer run stations. The largest price increase is at Arco, which increases its DTW's by $0.215 (0.267-0.052) per gallon on average. The one-price DTW for Arco ends up being at the max of its distribution of DTW's under price discrimination. Each of the branded refiners raises prices substantially to their dealers while simultaneously lowering prices significantly at their vertically integrated stations. The simulation shows that in markets where firms have franchises and vertically integrated outlets, banning wholesale price discrimination may effectively just move it to discrimination between vertically integrated and dealer stations. In addition, in a market with brand-loyalty, if upstream firms can't lower prices to 'meet competition' in markets

[25] Note that our simulation assumes that dealers follow the same mark-up behavior after the change in uniform pricing. It may be the case that wholesale price discrimination can be used by the refiner to enforce a competitive mark-up by the dealer. In this case, dealers may revert to Nash-Bertrand strategies, which would result in a larger rise in average prices under uniform wholesale pricing.

with less strongly branded competitors, they may opt to price completely to their brand-loyal segment only, thereby choosing a one-price that is close-to or at the highest-price.

The second column of Table 8 reinforces the first column. It shows how predicted volumes change under uniform wholesale pricing. Volumes sold fall at branded dealers and rise at the branded vertically integrated stations. Of all of the branded stations, Arco shifts the most volume from dealer-run stations to vertically integrated stations under uniform wholesale pricing. Over all stations, predicted total volume sold falls by 5.16% under uniform wholesale pricing, indicating an overall loss to welfare. Interestingly, prices at unbranded stations (the intercept in the regressions) remain essentially unchanged under uniform wholesale pricing, while volumes sold through these outlets increase significantly by a modest amount.

The results from the simulation show that knowing which stations paid higher DTWs and which stations paid lower ones under wholesale price discrimination would not necessarily help predict which markets would experience a rise in prices or a fall in prices under uniform wholesale pricing. Another way to see this is to regress the difference in predicted prices (uniform wholesale pricing – price discrimination) on the station characteristics, demographic characteristics, and measures of local competition used in the reduced-form analysis of wholesale price discrimination in Table 3. Table 8 presents these OLS results. Prices rise substantially at ARCO stations (relative to unbranded stations). Prices at Chevron, Shell and Unocal stations also rise, but by much less than they do at ARCO stations. In addition, prices rise at stations with service bays and at larger-format stations, fall at stations that are further away from the distribution rack. Importantly, prices rise most in high-poverty neighborhoods, increasing by three cents per gallon for a 10 percent increase in the poverty rate, suggesting that uniform wholesale pricing could be a regressive policy. The fact that many changes in prices are not as we would predict from the reduced-form results in Table 2 highlights the importance of taking into account market structure when conducting policy simulations. Solving for equilibrium prices using the first order conditions does depend heavily on functional form, but it includes in the predicted prices the optimal responses of other firms in the market given their characteristics and demand estimates, and the optimal strategy for firms who have vertically integrated stations to shift product to in response to regulation.

6.2 Full Vertical Integration

As demonstrated earlier, branded dealers appear to follow a constant mark-up policy over DTW cost. This may be because wholesale price discrimination is a sufficient tool for setting the retail price directly. Hastings (2004) finds no reduced-form evidence of price increases in markets that experience arguably exogenous increases in vertically integrated versus dealer-run stations, and suggests this is evidence that DTWs may be sufficient tools for setting retail prices directly, allowing the refiner to eliminate the double marginalization problem. We can examine this question in full given our structural model by simulating what prices would be if all stations were vertically integrated. This exercise does not just change a marginal station's ownership structure, but changes all branded stations at once to vertical integration. There may be equilibrium effects in this instance that may significantly impact prices. Firms playing Nash-bertrand game in differentiated products markets may find it profitable to delegate pricing decisions to an agent since that delegation acts to soften competition (Katz (1991), Rey and Stiglitz (1995)). Thus moving all firms to vertical integration may have a much larger impact than moving one marginal firm to vertical integration.

Figure 7 shows the results of the simulation. Average prices fall from \$1.793 under mixed-integration with wholesale price discrimination to \$1.779 under vertical integration. The density of stations in the lower-price range increases substantially. The difference between the two density plots suggests that refiners are not simply able to use DTW's to set retail prices as they would if they were able to set them directly. It may be that refiners still choose dealer-run stations because this contract type solves principal-agent problems associated with station management (Shepard (1993)), or it may be that they choose dealer-run formats to strategically soften competition. In the latter case, profits will actually be higher under the dealer-run format than under vertical integration.[26]

7. Conclusions

This paper estimated a structural model of retail gasoline markets in order to understand how wholesale price discrimination impacts differences in retail prices within markets, and what the impact of forcing uniform wholesale prices might be on retail prices levels and distribution. By examining actual wholesale price transactions data, we find that retailers, while they have local market power, set retail prices that are most consistent with a rule-of-thumb markup. They pass changes in wholesale costs quickly and completely through to retail, but keep a constant mark-up that varies across stations but remains constant over time as retail price levels and demand elasticities fluctuate. We also provide evidence that a random coefficients model where consumers display brand loyalty has stronger predictive power than a more typical model of vertical differentiation on brand quality.

Using our model estimates we simulate equilibrium prices under price discrimination and uniform wholesale pricing. We find that average prices would rise five cents per gallon under uniform wholesale pricing. We also find that while prices typically rise in areas where wholesale prices are lower under price discrimination, market structure is also a key factor in determining where prices would rise and where they would fall. Importantly, we predict that prices would rise in high-poverty areas and fall in high-income urban areas, demonstrating that policies that ban price discrimination can be regressive.

While these results are still preliminary, future versions will use the estimates to explore several other policy proposals regulating vertical contracts including divorcement and unbundled supply.

[26] Future versions of this paper will include an analysis that uses the simulation of profits to test between these two hypotheses (right now the simulations are still running).

REFERENCES

Asker, J., (2004), “Measuring Cost Advantages from Exclusive Dealing: An Empirical Study of Beer Distribution”. Working Paper, Harvard University.

Bajari, Patrick, J. Fox, K. Kim, and S. Ryan. (2008). “A Simple Nonparametric Estimator for the Distribution of Random Coefficients in Discrete Choice Models,” Unpublished manuscript, University of Minnesota.

Berry, Steven. (1994), “Estimating Discrete-Choice Models of Product Differentiation”, *RAND Journal of Economics*, 25 (2), 242–262.

Berry, Steven, M. Carnall and P. Spiller. (2006), “Airline Hubbing, Costs and Demand,” in Advances in Airline Economics, Vol. 1: Competition Policy and Anti-Trust, D. Lee, ed. Elsevier Press, pages 183-214.

Berry, Steven, James Levinsohn, and Ariel Pakes. (1995), “Automobile Prices in Market Equilibrium”, *Econometrica*, 63 (4), 841–890.

Berry, Steven, James Levinsohn, and Ariel Pakes. (2004), “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market,” *Journal of Political Economy*, 112 (1), 68-105.

Borenstein, S., Cameron, C, and Gilbert, R. (1997), “Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?”, *Quarterly Journal of Economics*, 112 (1), 306–339.

Borenstein, S. and A. Shepard. (2002), “Sticky Prices, Inventories, and Market Power in Wholesale Gasoline Markets”, *RAND Journal of Economics*, 33 (1), 116–139.

Briesch, Richard, P. Chintagunta and R. Matzkin; (2002) "Semiparametric Estimation of Brand Choice Behavior." *Journal of the American Statistical Association*, *97*(460), pp. 973-82.

Chandra, Ambarish and M. Tappata. (2008) “Consumer Search and Dynamic Price Dispersion: An Application to Gasoline Markets.” Working Paper, University of British Columbia.

Chevalier, Judith, D. Harrington and F. Soctt-Morton, (2008). “Differentiated to Death?” Manuscript, Yale University.

Chintagunta, Pradeep, E. Kyriazidou and J. Perktold. (2001), "Panel Data Analysis of Household Brand Choices." *Journal of Econometrics*, 2001, *103*(1-2), pp. 111-53.

Chintagunta, Pradeep, D. Jain and N. Vilcassim, (1991), "Investigating Heterogeneity in Brand Preferences in Logit Models for Panel Data." *Journal of Marketing Research*, *28*(4), pp. 417-28.

Chintagunta, Pradeep. (1992) "Estimating a Multinomial Probit Model of Brand Choice Using the Method of Simulated Moments." *Marketing Science*, 1992, *11*(4), pp. 386-407.

Corts, Kenneth S. (1998). “Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment,” *RAND Journal of Economics*, 29(2), pp. 306-323.

Davis, Peter (2006). ‘‘Spatial Competition in Retail Markets: movie theaters,’’ *The RAND Journal of Economics* 37(4), pp. 964–82.

Hastings, Justine. (2004), “Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California,” *American Economic Review*, 94(1), 317-328.

Hastings, Justine and R. Gilbert. (2005), “Vertical Integration in Gasoline Supply: An Empirical Test of Raising Rivals' Costs.”, *Journal of Industrial Economics*, 53 (4), 437-571.

Hastings, Justine, T. Kane and D. Staiger. (2007), “Preferences and Heterogeneous Treatment Effects in a Public School Choice Lottery,” National Bureau of Economic Research Paper Working Paper No. 12145.

Hellerstein, R. (2008), “Who Bears the Cost of a Change in the Exchange Rate? Pass-through Accounting for the Case of Beer,” *The Journal of International Economics,* forthcoming.

Hendricks, K. and McAfee, R. P., (2006), “A Theory of Bilateral Oligopoly,” Working Paper, University of Texas at Austin Department of Economics.

Hortaçsu, Ali and C. Syverson (2007), “Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices,” *Journal of Political Economy*, 115 (2), 250-301.

Houde, J.F. (2008), “Spatial Differentiation in Retail Markets for Gasoline,” Working Paper, University of Wisconsin at Madison.

Jain, Dipak, N. Vilcassim and P. Chintagunta, (1994), "A Random-Coefficients Logit Brand-Choice Model Applied to Panel Data." *Journal of Business and Economic Statistics*, 12(3), pp. 317-28.

Kalouptsidi, Myrto. (2008), "From Market Shares to Consumer Types: Duality in Differentiated Product Demand Estimation" Unpublished manuscript, Yale University.

Katz, M. (1987), "The Welfare Effects of Third-Degree Price Discrimination in Intermediate Goods Markets," *American Economic Review*, 77, 154-67.

Michael L. Katz, (1991), "Game-Playing Agents: Unobservable Contracts as Precommitments," *RAND Journal of Economics*, 22(3), 307-328.

Lewis, Matthew. (2005). "Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market." Working paper, Ohio State University.

Lewis, Matthew and H. Marvel. (2007), "When Do Consumers Search?" Working paper, Ohio State University.

Lewis, Matthew. (2008). "Price Dispersion and Competition with Differentiated Sellers," *Journal of Industrial Economics*, 56 (3), pp. 654-678.

McAffee, R. P. and M. Schwartz. (1994), "Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity and Uniformity," *American Economic Review*, 84, 210-30.

Mortimer, J. H. (2008). "Vertical Contracts in the Video Rental Industry," *The Review of Economic Studies*, forthcoming.

Nevo, Aviv. (2000), "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," *The RAND Journal of Economics*, 31(3), 395-421, 2000.

Nevo, Aviv. (2000), "A Practitioner's Guide to Estimation of Random Coefficients Logit Models of Demand," *Journal of Economics & Management Strategy*, 9(4), 513-548.

Nevo, Aviv. (2001), "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69(2), 307-342.

Nevo, A. and F. Rossi. (2008), "An Approach for Extending Dynamic Models to Settings with Multi-Product Firms," *Economics Letters*, forthcoming.

Noel, Michael D. (2007a), "Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market," *Journal of Industrial Economics*, 55 (1), 69-92.

Noel, Michael D. (2007b), "Edgeworth Price Cycles, Cost-based Pricing, and Sticky Pricing in Retail Gasoline Markets," *Review of Economics and Statistics*, 89 (2), 324-334.

Ordover, J. A.; Saloner, G. and Salop, S., (1990), “Equilibrium Vertical Foreclosure,” *American Economic* Review, 80, 127–142.

Rey, Patrick and J. Stiglitz, (1995), “The role of exclusive territories in producers’ competition,” *RAND Journal of Economics*, 26 (3), 431-51.

Salop, S. C. and Scheffman, D. T. (1987), “Cost-Raising Strategies,” *Journal of Industrial Economics*, 36 (1), 19–34.

Shepard, Andrea, (1993) "Contractual Form, Retail Price, and Asset Characteristics." *RAND Journal of Economics*, 24(1), pp. 58-77.

Slade, M. E. (1992), “Vancouver’s Gasoline-Price Wars: An Empirical Exercise in Uncovering Supergame Strategies”, *Review of Economic Studies*, 59 (2), 257-76.

Sorensen, Alan. (2000), "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs," *Journal of Political Economy*, 108 (4), 833-62.

Sriram, S., P. Chintagunta and R. Neelamegham. (2006), "Effects of Brand Preference, Product Attributes, and Marketing Mix Variables in Technology Product Markets." *Marketing Science*, 25(5), pp. 440.

Verlinda, Jeremy. (2007) “Price-Response Asymmetry and Spatial Differentiation in Local Retail Gasoline Markets,” *DOJ EAG 07-4*, Discussion paper.

Villas-Boas, S. B., (2007). “Vertical Relationships between Manufacturers and Retailers: Inference With Limited Data,” *The Review of Economics Studies*, 74 (2), 625-652.

Villas-Boas, S. B., (2009). “An Empirical Investigation of the Welfare Effects of Banning Wholesale Price Discrimination,” *RAND Journal of Economics*, *forthcoming*.

Villas-Boas, S. B. and Hellerstein, R. (2006), “Identification of Supply Models of Retailer and Manufacturer Oligopoly Pricing,” *Economics Letters*, 90 (1), 132–140.

Yoshida, Y. (2000), “Third-Degree Price Discrimination in Input Markets: Output and Welfare,” *American Economic Review*, 90, 240-246.

Figure 1: Average Weekly CA Retail, LA Spot, NY Spot and Crude Oil Prices (Excise Tax added to Spot & Crude for Comparison.)

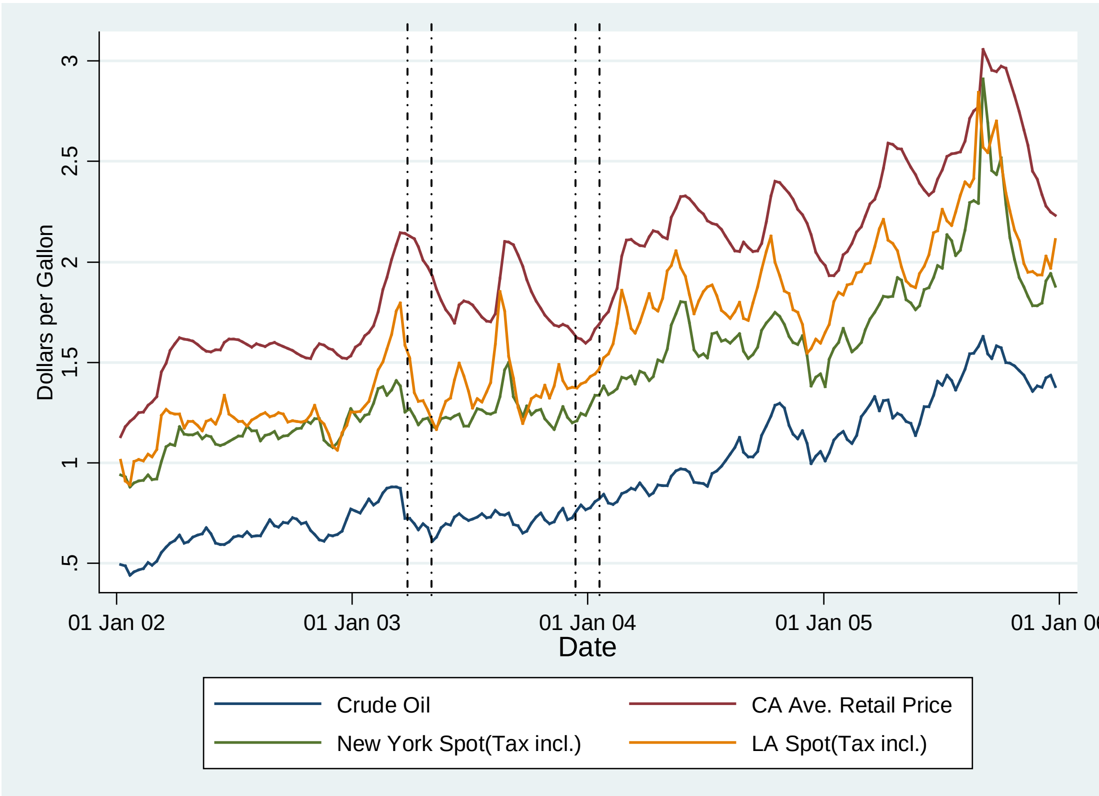

Notes: Weekly data on Crude oil prices, New York Spot prices and California Average Retail Prices were downloaded from the Energy Information Administration website. Los Angeles spot price is from Oil Price Information Service and are the weekly average of the average daily transaction price. Excise taxes of $0.348 cents were added to spot prices and crude oil prices for easier visual comparison with retail price variation.

Figure 2: Mean Retail (ex-tax), Mean DTW, Mean Spot Price

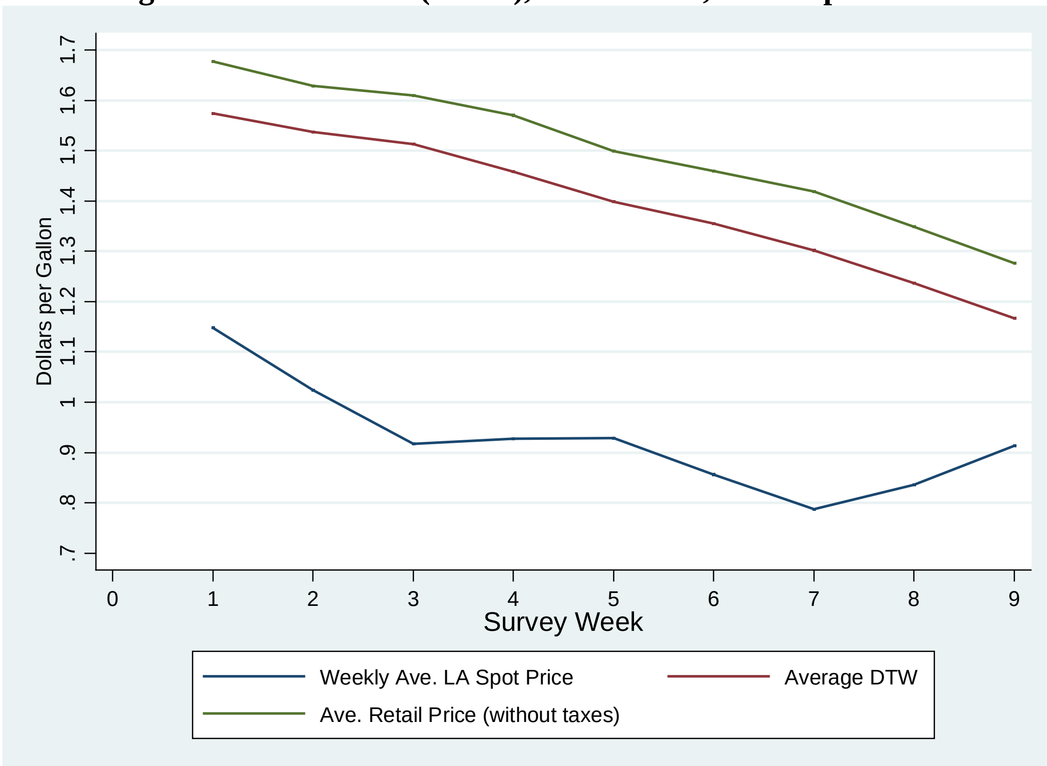

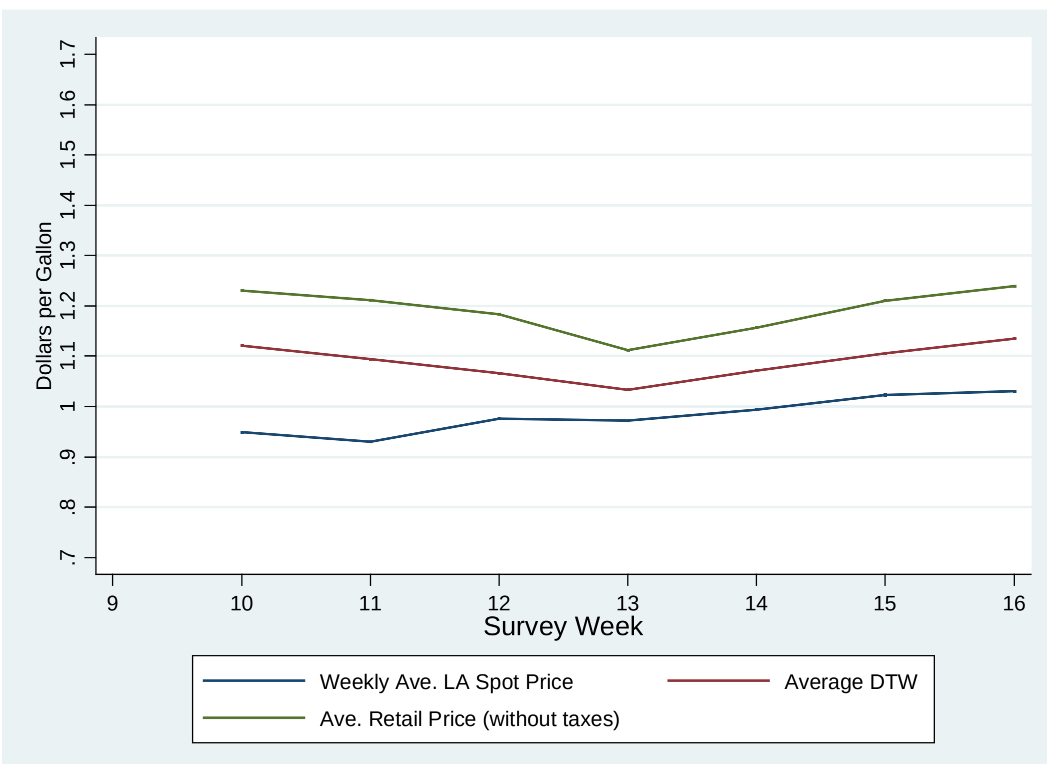

Notes: Retail and DTW data reported by survey participants. Los Angeles Spot Price comes from Oil Price Information Service. Retail prices exclude 7.75% sales tax and $0.348 total excised taxes.

Figure 3: Margins for Refiners and Retailers; Means, 90th, 10th percentiles

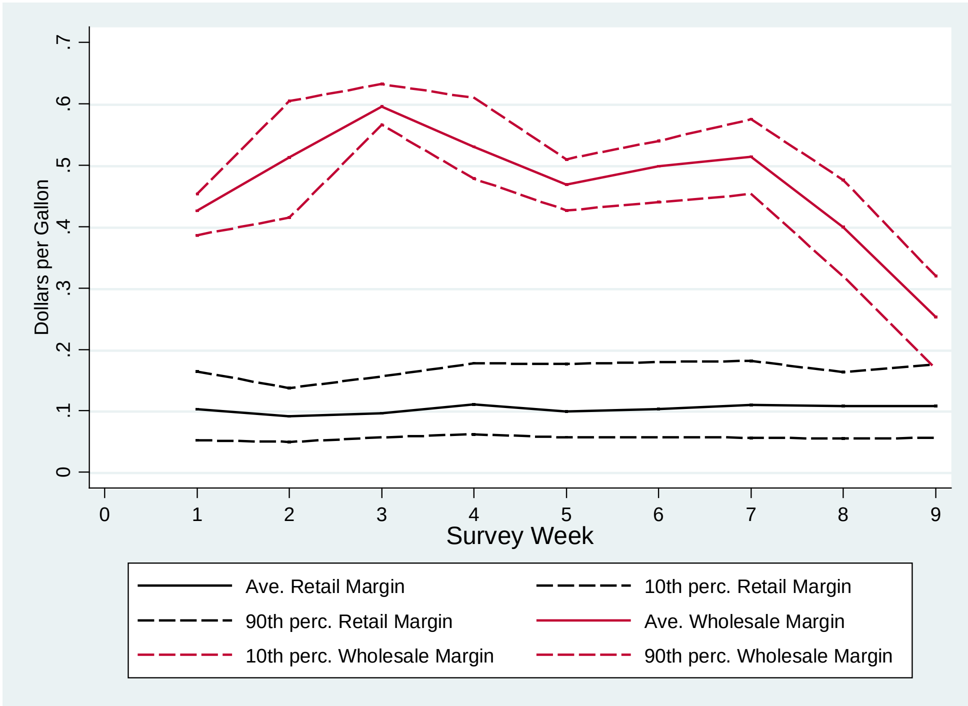

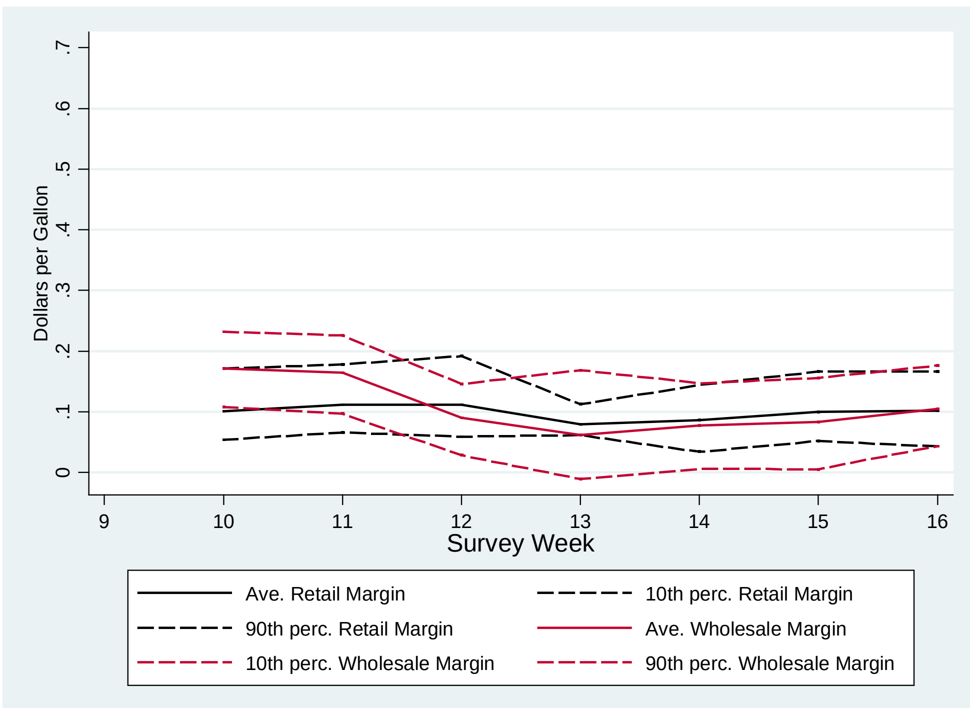

Figure 4: Map of San Diego Stations

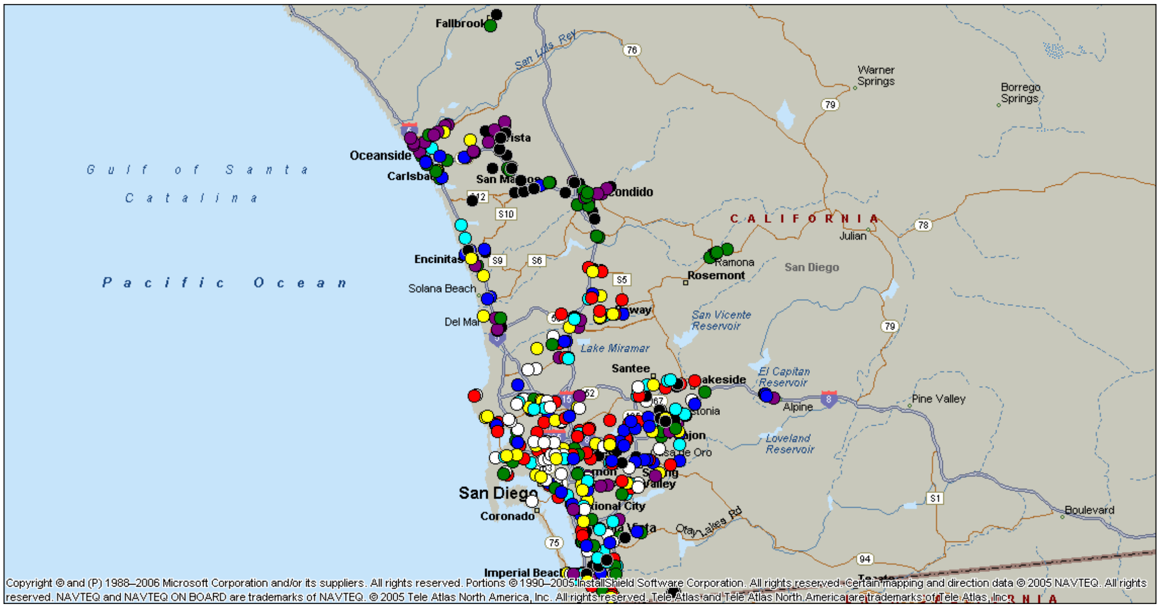

Figure 5: Distribution of Prediction Errors by Model

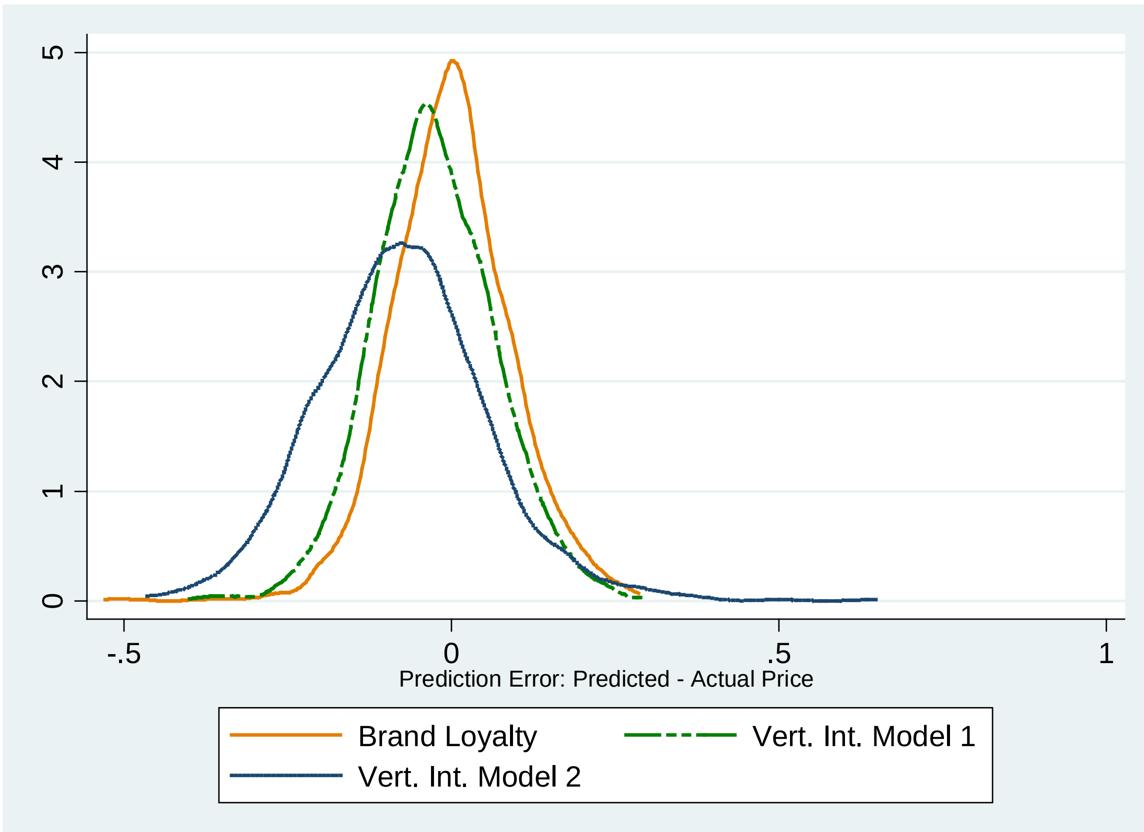

Figure 6: Distribution of Prices Under Price Discrimination vs. Uniform Wholesale Prices

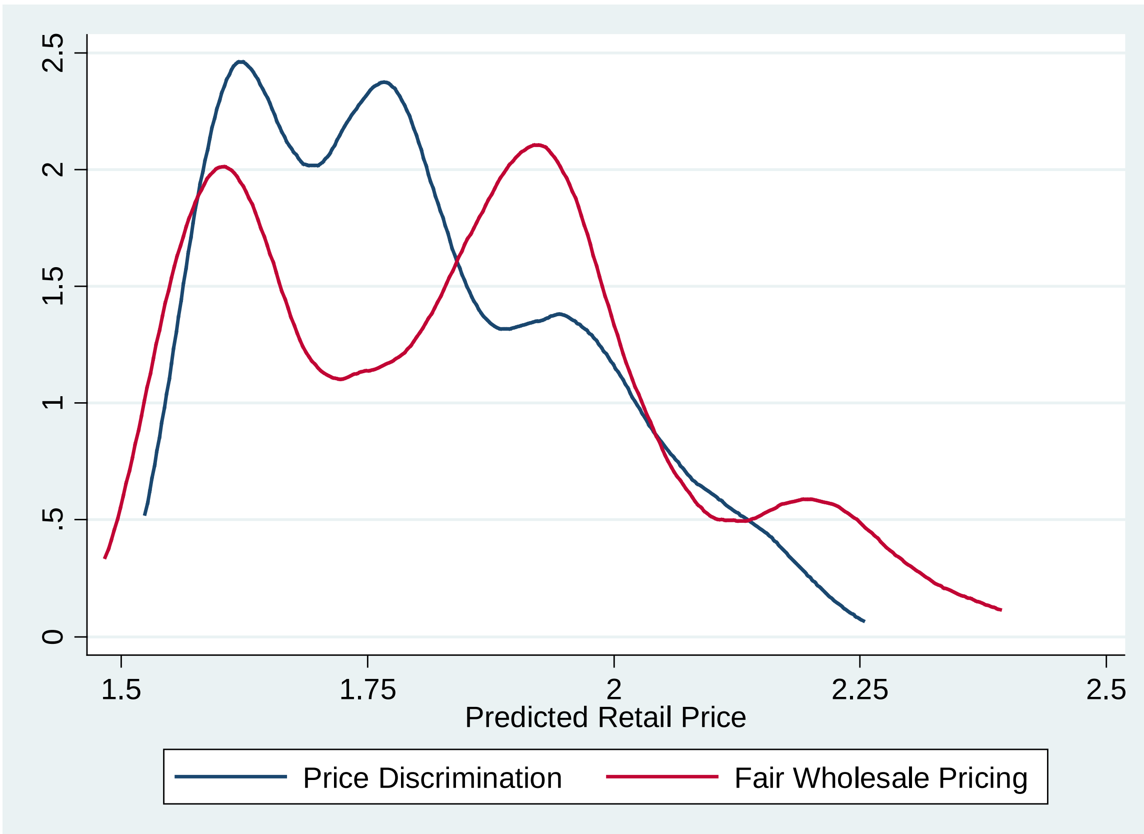

Figure 7: Distribution of Prices Under Full Vertical Integration versus Price Discrimination with Dealer-run Stations with Rule-of-Thumb Mark-ups

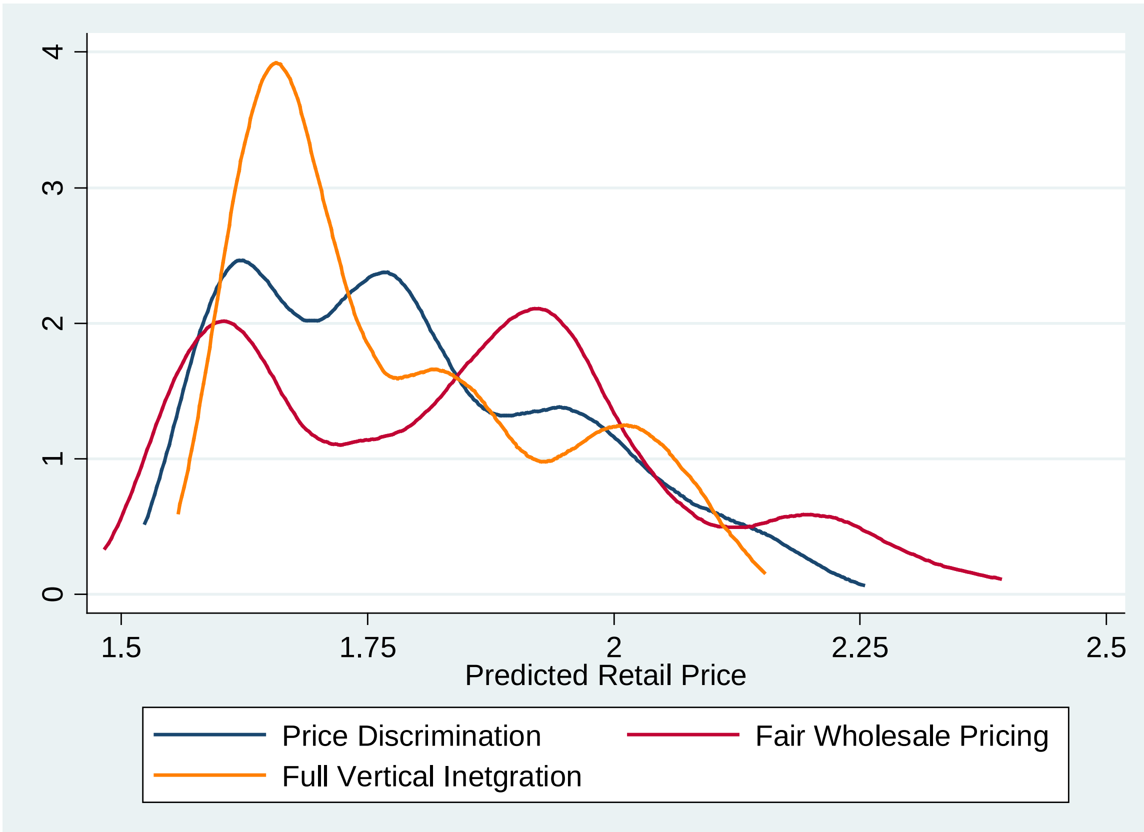

Table 1: Characteristics of Stations in Samples

	All Stations	Price and Volume Sample	All Contract Dealer stations	DTW Survey Participants
	Mean	Mean	Mean	Mean
Station Characteristics				
ARCO	0.204	0.216	0.209	0.217
Chevron	0.122	0.124	0.159	0.117
Mobil/Exxon	0.136	0.117	0.137	0.142
Shell/Texaco	0.154	0.168	0.303	0.317
Unocal	0.117	0.124	0.191	0.192
7-Eleven	0.099	0.096	--	--
Other	0.164	0.148	--	--
Fueling Positions	8.281	8.258	9.129	8.967
Car Wash	0.152	0.152	0.184	0.256
Local Competition				
N Stations in 1 mile**	4.000	4.004	3.872	3.578
HHI of Stations in 1 mile**	4,136	4,125	3,926	4,124
*Local Demographics****				
Average Income	46,147	46,094	47,122	46,362
Fraction White	0.668	0.667	0.678	0.687
N Children per Household	0.944	0.947	0.904	0.916
Average Commute Time	25.01	25.075	24.416	25.047
N. Cars Per Household	1.679	1.681	1.668	1.688
N	609	547	277	120

Notes: *Denotes a sample mean that is statistically significantly different than the population mean given the sample size and population variance. All sample statistics are consistent with a random sample from the population. **HHI defined as the share of total fueling positions a retailer has among stations within a one mile driving distance. Driving distance is defined using mapping and routing software that calculates routes using a street network. ***Local demographics are taken from 2000 Census variables at the block group level, using a population-weighted average of statistics for all block groups with centroids within 1.5 miles of a station.

Table 2: Estimates of DTW to Retail Pass-through Rates

Dependent Var: Retail price ex-tax)	(1) OLS	(2) IV
Dealer Tank-wagon Price	0.983**	0.992**
	(0.014)	(0.019)
Station Fixed Effects	Y	Y
		Excluded Instruments: LA spot price*season LA spot price*distance to rack LA spot price* brand dummies
Observations	1462	1462
R-squared	0.97	

Notes: *significant at 5%; ** significant at 1%. Standard errors in parentheses are clustered at the station level.

Table 3: Factors that Influence Wholesale versus Retail Prices

	(1) Wholesale Price		(2) Retail Margin	
Station Characteristics				
Chevron	0.057**	(0.006)	0.029	(0.015)
Mobil/Exxon	0.074**	(0.006)	-0.009	(0.017)
Shell/Texaco	0.036**	(0.006)	0.002	(0.016)
Union 76	0.068**	(0.007)	0.027	(0.018)
Located on a freeway exit	0.011**	(0.004)	0.022*	(0.009)
Distance to distribution rack (miles)	0.001**	(0.000)	-0.001	(0.001)
Station has a car wash	-0.005	(0.006)	0.028	(0.019)
Station has a service bay	0.003	(0.005)	-0.008	(0.013)
Station has >600 sqft of convenience store space	0.003	(0.005)	-0.022	(0.015)
Station offers full service	0.013	(0.009)	0.041*	(0.020)
Number of fueling positions	-0.001	(0.001)	-0.003	(0.002)
Local Demographics				
Block group rental rate	-0.0002	(0.003)	-0.003*	(0.001)
Neighborhood median income	0.011**	(0.003)	-0.002	(0.006)
Neighborhood % on public assistance	0.009	(0.115)	-0.135	(0.246)
Neighborhood % white	-0.040*	(0.017)	-0.019	(0.043)
Neighborhood # of kids/household	-0.003	(0.013)	-0.030	(0.032)
Neighborhood average commute time (minutes)	-0.005**	(0.001)	-0.003	(0.003)
Neighborhood average number of cars per household	-0.041**	(0.011)	0.040	(0.029)
Local Competition				
HHI of stations within 1 mile	0.000	(0.002)	0.003	(0.004)
Number stations within 1 mile	0.003	(0.002)	0.003	(0.006)
Number of stations of same brand within 1 mile	-0.011**	(0.004)	-0.017	(0.010)
Number of unbranded stations within 1 mile	-0.005*	(0.002)	-0.008	(0.007)
Number of refiner owned and operated stations within 1 mile	-0.008	(0.006)	-0.031	(0.017)
Average driving distance to stations within 1 mile	-0.002	(0.007)	-0.005	(0.020)
Date Fixed Effects	Y		Y	
Observations	1206		1162	
R-squared	0.98		0.34	

Notes: Standard errors in parentheses. * significant at 5%; ** significant at 1%.

Table 4: Demand Parameters and Elasticities for Vertical Differentiation Model 1

	Specification 1	
	Coefficient	Standard Error
Mean preferences		
Price	-2.578*	(0.043)
Station has a carwash	0.163*	(0.064)
Station has a service bay	-0.284*	(0.052)
Station offers full service	-0.065	(0.089)
Convenience Store>600sqft	0.032	(0.048)
Number of fueling positions	0.025*	(0.008)
Station located on a highway	0.295*	(0.040)
7-Eleven	-0.827*	(0.083)
Arco	0.151*	(0.064)
Chevron	0.336*	(0.071)
Mobil	-0.088	(0.071)
Shell	1.167*	(0.074)
Unocal	0.244*	(0.071)
Interaction with Income		
Arco	0.015*	(0.004)
Chevron	-0.038	(0.048)
Mobil	-0.004	(0.003)
Shell	-0.149*	(0.004)
Unocal	-0.057*	(0.027)
Price	0.090*	(0.001)
St. Dev. Idiosyncratic Taste		
Arco	0.058	(0.419)
Chevron	0.075	(11.663)
Mobil	0.203	(0.664)
Shell	0.006	(0.823)
Unocal	0.070	(6.743)
Price	0.037	(0.128)
*Demand Elasticities***	*Mean*	*St. Dev.*
Elasticity Period 1	-3.535	0.512
Elasticity Period 2	-3.306	0.449
Elasticity Period 3	-3.015	0.385
Elasticity Period 4	-2.989	0.386

Notes: * significant at 5% level. Parameter estimates from GMM estimation. **All demand elasticities are significantly different from zero at the one percent level or higher.

Table 5: Demand Parameters and Elasticities for Vertical Differentiation Model 2

	Specification 2	
	Coefficient	Standard Error
Mean preferences		
Price	-2.807*	(0.444)
Station has a carwash	0.159*	(0.067)
Station has a service bay	-0.309*	(0.054)
Station offers full service	-0.047	(0.093)
Convenience Store>600sqft	-0.005	(0.049)
Number of fueling positions	0.030*	(0.008)
Station located on a highway	0.267*	(0.042)
7-Eleven	-0.820*	(0.086)
Arco	0.440*	(0.066)
Chevron	0.123*	(0.073)
Mobil	-0.034	(0.073)
Shell	2.565*	(0.076)
Unocal	-0.251*	(0.073)
Interaction with Income		
Arco	-0.008	0.066
Chevron	-0.166*	0.066
Mobil	-0.094	0.060
Shell	-0.281*	0.085
Unocal	-0.127*	0.060
Price	0.168*	0.023
Interaction with Commute Time		
Arco	0.006	0.071
Chevron	0.053	0.050
Mobil	0.035	0.056
Shell	-0.012	0.091
Unocal	0.046	0.051
Price	-0.016	0.021
St. Dev. Idiosyncratic Taste		
Arco	0.261	(0.310)
Chevron	0.085	(0.798)
Mobil	0.027	(1.079)
Shell	0.020	(0.792)
Unocal	0.022	(0.968)
Price	0.169*	(0.048)
*Demand Elasticities***	*Mean*	*St. Dev.*
Elasticity Period 1	-3.363	0.462
Elasticity Period 2	-3.141	0.403
Elasticity Period 3	-2.848	0.348
Elasticity Period 4	-2.822	0.346

Notes: * significant at 5% level. Parameter estimates from GMM estimation. **Most of the demand elastcities are insignificant at the 5% level.

Table 6: Demand Parameters and Elasticities for Brand Loyalty Model

	Specification 3	
	Coefficient	St. Error
Mean preferences		
Price	-2.661*	(0.0395)
Station has a carwash	0.137*	(0.0569)
Station has a service bay	-0.309*	(0.0506)
Station offers full service	-0.116	(0.0882)
Convenience Store>600sqft	0.057	(0.0439)
Number of fueling positions	0.026*	(0.0075)
Station located on a highway	0.288*	(0.0402)
7-Eleven	-0.770*	(0.0747)
Brand-Loyal Type Share		
Arco	0.199*	(0.0003)
Chevron	0.235*	(0.0641)
Mobil	0.136*	(0.0044)
Shell	0.141*	(0.0648)
Unocal	0.037	(0.1097)
Implied share of "Shoppers"	*0.252*	
Brand-Loyal Type Value		
Arco	1.041*	(0.2329)
Chevron	0.737*	(0.0083)
Mobil	0.077	(0.0766)
Shell	1.123*	(0.1865)
Unocal	0.154*	(0.0067)
Idiosyncratic Preferences for Price		
Income	0.075*	(0.0041)
Commute Time	0.006*	(0.0027)
St. Dev. of Normal	0.020*	(0.0058)
*Demand Elasticities***	*Mean*	*St. Dev.*
Elasticity Period 1	-3.691	0.279
Elasticity Period 2	-3.408	0.278
Elasticity Period 3	-3.070	0.235
Elasticity Period 4	-3.043	0.236

Notes: * significant at 5% level. Parameter estimates from GMM estimation. **All demand elasticities are significantly different from zero at the one percent level or higher.

Table 7: Estimating Parameters of Marginal Cost

	Brand Loyalty	Vert.Int. Model 1	Vert.Int. Model 2
Spot Price	0.991*	1.003*	1.054*
	(0.076)	(0.087)	(0.091)
Station has car wash	0.002	0.040*	0.041*
	(0.006)	(0.007)	(0.008)
Station has service bay	0.034*	0.041*	0.039*
	(0.006)	(0.007)	(0.007)
Station has full service	-0.040*	-0.027*	-0.054*
	(0.010)	(0.012)	(0.012)
Convenience Store > 600sqft	-0.028*	-0.036*	-0.023*
	(0.005)	(0.006)	(0.006)
Distance to distribution rack	-0.001*	-0.001*	-0.002*
	(0.000)	(0.000)	(0.000)
Constant	0.681*	0.645*	0.567*
	(0.065)	(0.074)	(0.077)
Time dummy 2	-0.164*	-0.161*	-0.158*
	(0.007)	(0.008)	(0.008)
Time dummy 3	-0.370*	-0.368*	-0.367*
	(0.006)	(0.007)	(0.008)
Time dummy 4	-0.450*	-0.449*	-0.452*
	(0.009)	(0.010)	1.054
Observations	1576	1576	1576

Notes: Standard errors in parentheses, * significant at 5%.

Table 8: Changes in Prices and Volumes by Contract and Brand

	Predicted Price Difference (Uniform – Price Discrim.) (1)	Predicted Volume Difference (Uniform – Price Discrim.) (2)
Arco	-0.052**	19,510.308**
	(0.009)	(3,548.970)
Chevron	-0.026**	10,353.213**
	(0.008)	(2,958.513)
Mobil	-0.028**	7,723.592**
	(0.005)	(1,564.559)
Shell	-0.032**	13,848.337**
	(0.008)	(3,438.870)
Unocal	-0.023**	4,599.400**
	(0.007)	(1,440.423)
Arco Dealer	0.267**	-64,783.050**
	(0.013)	(4,844.073)
Chevron Dealer	0.096**	-28,079.271**
	(0.011)	(3,938.083)
Mobil Dealer	0.050**	-10,353.438**
	(0.008)	(1,900.796)
Shell Dealer	0.142**	-36,924.275**
	(0.011)	(3,980.241)
Unocal Dealer	0.132**	-21,836.478**
	(0.009)	(1,772.183)
Constant (unbranded)	-0.001**	1,882.775**
	0.000	(130.353)
Observations	1,576	1,576
R-squared	0.77	0.62

Note: * significant at 5%; ** significant at 1%. Standard errors clustered at the station level. Robust standard errors in parentheses.

Table 9: Where Do Prices Change Under Uniform Price Policy?

Dep. Var: Difference in Predicted Prices (Uniform – Price Discrimination)	Coefficient	St. Error
Station Characteristics		
ARCO	0.122**	(0.007)
Chevron	0.025**	(0.008)
Mobil/Exxon	-0.005	(0.008)
Shell/Texaco	0.078**	(0.008)
Union 76	0.067**	(0.008)
Located on a freeway exit	0.005	(0.009)
Distance to distribution rack (miles)	-0.019**	(0.005)
Distance to Distribution Rack	0.000	0.000
Station has a car wash	-0.005	(0.007)
Station has a service bay	0.031**	(0.006)
Convenience store >600 sqft	0.003	(0.005)
Station offers full service	-0.012	(0.010)
Number of fueling positions	0.005**	(0.001)
Local Demographics		
Block group rental rate	0.000	(0.001)
Neighborhood median income	0.000	0.000
Neighborhood % on public assistance	0.304*	(0.135)
Neighborhood % white	0.014	(0.020)
Neighborhood # of kids/household	-0.018	(0.017)
Neighborhood average commute time (minutes)	0.000	(0.001)
Neighborhood average number of cars per household	0.026*	(0.014)
Local Measures of Competition		
HHI of stations within 1 mile	0.000*	0.000
Number stations within 1 mile	0.002	(0.002)
Number of stations of same brand within 1 mile	-0.032**	(0.004)
Number of unbranded stations within 1 mile	0.001	(0.003)
Number of refiner owned and operated stations within 1 mile	0.021**	(0.007)
Observations	1,576	
R-squared	0.300	

Notes: Standard errors in parentheses, * significant at 5%; ** significant at 1%.