Abstract

This study explores the possibility that local market power influences the observed asymmetric relationship between changes in wholesale gasoline costs and changes in retail gasoline prices. I exploit an original data set of weekly gas station prices in Southern California from September 2002 to May 2003, and take advantage of detailed station and local market level characteristics to determine the extent to which spatial differentiation influences price response asymmetry. I find that brand identity, proximity to rival stations, bundling and advertising, operation type, and local market features and demographics each influence a station's predicted price-response asymmetry.

1 Introduction

The pricing dynamics of retail gasoline have resurfaced as an important policy issue in the last few years. Of particular concern is the observed phenomenon that prices respond asymmetrically to cost shocks—prices rise at a much faster rate with cost increases than they do with cost declines. Research into the possible sources of this asymmetry remains an ongoing concern. In this study, I investigate the extent to which local market power may contribute to price asymmetry.

Past studies have found only limited support for the hypothesis that market power influences pricing dynamics. Part of the difficulty is that for the retail sector in the gasoline industry, market power occurs at the local or station level through geographic and other forms of spatial differentiation. However, most gasoline pricing studies employ publicly available aggregate data on prices, usually at a regional level. The few emerging studies that take advantage of station-level data are generally cross-sections, which are unable by definition to describe the types of pricing dynamics discussed here.

I overcome these data complications by collecting a station-level data set of weekly retail gasoline prices from September 2002 to May 2003, which I use to examine how station-specific characteristics such as location relative to competitors, consumers, and major infrastructures are potential sources of market power, and thus how they contribute to pricing asymmetry at a micro level. In this paper, I establish that price-response asymmetry is a dominant feature in my data set, and explore the possibility that certain station-level features such as brand identity and other site- and local-market characteristics exhibit varying degrees of asymmetry consistent with a positive local market power effect. I assume that each station has its own response to current and lagged cost and price changes, that these effects depend explicitly on its site- and local-market specific characteristics, and that these effects are themselves correlated with each other as a function of the geographic distance between stations.

Specifically, I find that, aggregated across all stations, three weeks after a singleton wholesale cost increase of 100¢, retail prices are predicted to rise an estimated 110¢; but when costs fall by 100¢, retail prices only fall by 83¢. This 27¢ difference, which I define as the price-response asymmetry, declines gradually toward zero after the third week until retail prices settle at their estimated long-run response. I also find that brand identity contributes measurably to asymmetry: three weeks after a wholesale cost shock, the asymmetry for the highest-priced brands is estimated at 34¢, while for unbranded stations the difference in responses is estimated at only 14¢. Regarding the benefit of geographic isolation, I find that stations with no rivals in immediate proximity exhibit an asymmetry after the third week that is approximately 7¢ greater than if these same stations had a neighbor immediately nearby. I also find that after the third week following a cost shock, the difference in asymmetry for stations with greater versus lesser ease-of-access is approximately 20¢, diminishing slowly toward zero afterward. Similar results, though different in magnitude, are attributed to stations with a convenience store and to stations with relatively more pumps on their lot. Lastly, I find that more competitive local markets have a lower price-response asymmetry by 10¢ at the third week after a cost shock, while a one standard deviation increase in the local market population size is associated with a 5¢ increase in price-response asymmetry.

The remainder of this paper is organized as follows. In Section 2, I suggest a theoretical link between market power and pricing asymmetry, and some of the existing empirical and theoretical literature that supports it. In Section 3, I describe the basic empirical model and demonstrate the existence of pricing asymmetry in my data set. I follow this up in Section 4 by introducing an empirical model that allows me to predict price-response asymmetry for varying levels of the station characteristics. I subsequently estimate the model and discuss the results and their implications for drawing a conclusion on the link between local market power and price-response asymmetry. Concluding remarks and suggestions for future research are offered in Section 5.

2 Market Power as a Source of Price-Response Asymmetry

California consumers witnessed considerable volatility in retail gasoline prices in 2003, with price spikes that surpassed those seen in the rest of the U.S.. Current industry analysis attributes California's price volatility to industry responses to a series of local supply disruptions in a market where refiners typically operate near full capacity with limited capability to temporarily expand production. Figure 1 presents a chart of the retail prices for the Los Angeles basin, along with corresponding spot prices for Los Angeles reformulated gasoline.[1]

In addition to attributing price volatility in 2003 to an unusually high frequency of supply disruptions, Figure 1 also demonstrates at a highly visible level the presence of price-response asymmetry during this time period. During a supply shock, spot prices rise quickly to clear the wholesale market for gasoline. Once the supply disruption is alleviated, spot prices fall accordingly, though not quite so fast as they rose. Retail prices, however, respond much more slowly to decreases in wholesale costs than increases; the reaction to cost increases appears nearly immediate, while the reaction to cost decreases appears to take several weeks.

In this study, I explore the possibility that the dynamic phenomena observed in gasoline prices are influenced by local market power.[2] In separate but related literatures, empirical studies of the gasoline industry find: (1) that stations do enjoy local market power that derives from their ability to spatially differentiate their stores from potential competitors,[3] and (2) that station-level pricing data are consistent with certain models of tacitly collusive

[1] Source data for the chart are from the Energy Information Administration (EIA). The EIA also issued a public report (EIA, 2003) at the request of the U.S. House Subcomittee on Energy Policy, detailing the sources of the California supply shocks and their effects on local prices. The identified shocks are listed in Figure 1 where appropriate.

[2] Throughout this paper, market power is formally defined as the ability of a firm to maintain its customer base when the price differential between comparable own and rival goods increases.

[3] The seminal papers that establish local, retail-level market power in in the gasoline industry are Borenstein (1991) and Shepard (1991). Structural evidence of market power has also surfaced in the retail gasoline literature recently, notably in Manuszak (2001), and Romley (2002). Hastings (2000) and Barron et al. (2003) exploit natural or designed experiments on station-level pricing, also finding strong evidence of local, station-level, market power.

behavior.[4] These empirical findings would seem to corroborate the hypothesis that price-response asymmetry is amplified by the presence of stations with market power engaging in tacit collusion with their neighbors. If true, we would expect to find that variations in local market power are positively correlated with fluctuations in the measured asymmetry of price responses to cost shocks.

The theoretical model that I use to link local market power to price-response asymmetry is essentially built on the discussion in Borenstein and Shepard (1996), which offers a compelling summary of the theoretical argument for upstream cost-oriented asymmetry in prices (markups). Rotemberg and Saloner (1986; hereafter RS), demonstrate a model of tacit collusion that leads to price wars during unexpected demand shocks. Borenstein and Shepard reinterpret the Haltiwanger and Harrington (1991; hereafter HH) refinement of the RS model, arguing that expectations of future cost changes result in asymmetric price-responses depending on the direction of the cost change.

The underlying setup of these models is that firms play a repeated Bertrand game with a grim trigger strategy. All firms face equivalent market demand and constant marginal costs. In the absense of collusion, all firms charge marginal cost and earn zero economic rents. In a perfectly collusive outcome, each firm charges the monopoly price and receives $1/n$ share of the monopoly economic rent. Under the grim trigger strategy, firms will charge the collusive price so long as no firm deviates, but after a firm deviates, the other firms respond by reverting to static Bertrand competition. Although the deviating firm receives a one-period gain by stealing the market, in future periods he earns zero profits. Generally, provided the common discount factor is greater than $(n-1)/n$, no firm will ever find it profitable to deviate and so all firms earn a positive markup in all periods.

RS modify the basic supergame model of tacit collusion by introducing stochastic, *iid* demand shocks, the main affect of which is that price wars occur when demand is high

[4] Slade (1992) finds evidence in the price war behavior of retail gasoline stations in Vancouver, BC that is consistent with a kinked demand model of tacit collusion. Other models of tacit collusion have also been considered for the retail gasoline industry. In particular, Borenstein and Shepard (1996) use a city-level panel data set to show that US prices are consistent with a variant of Rotemberg and Saloner's (1986) model of tacit collusion.

because in the near term collusive profits are less than they would otherwise be.[5] The reduction in collusive profits increases the incentive for any one firm to deviate, and so the collusive price must temporarily fall to remove that incentive. The reverse is true when demand is low.

HH further modify this basic model. Instead of stochastic *iid* demand shocks, they imagine that the dynamic path of demand is deterministic (but changing). This leads to a model where collusive prices and the ease of collusion depend not only on the level of demand but also on the near-term future changes in demand. The main directional effect in HH is that, conditioning on the level of demand, collusive prices are increasing when demand is increasing and decreasing when demand is decreasing. As with RS, this occurs because of the effect of near-term demand changes on the near-term profits from collusion. When demand is increasing, near-term collusive profits increase, making collusion relatively more sustainable.

Borenstein and Shepard (1996) note that the RS and HH models are trivially recast in terms of dynamic changes in marginal cost rather than demand. It follows from the RS model that if firms observe a high-cost state, then near-term future collusive profits are relatively higher because costs are expected to fall in the next period. This makes collusion easier to sustain and leads to higher collusive markups. Conversely, when firms observe a low-cost state, collusive markups will fall because costs are expected to increase in the next period. This logic follows into the HH model: markups should be increasing when future costs are decreasing, and markups should be decreasing when future costs are increasing. In each of these models, we obtain a result where prices respond asymmetrically to cost changes.

Both RS and HH identify the same underlying dependence between the number of firms in the market and the common discount factor that exists in the basic supergame model of

[5] The statement "demand is high" in the RS model means that firms observe a current period demand shock that exceeds μ, the expected value of demand. Since demand is *iid*, the following period's expectation of demand is μ, and so each firm's expectation of near-term collusive profits is lower relative to the potential gains from cheating today. This is because each firm expects tomorrow's demand shock to be lower than today's realization of demand.

tacit collusion: collusion is assured if the discount factor δ is such that $\delta \geq (n-1)/n = \overline{\delta}$, and perfect collusion (monopoly pricing) is assured when $\delta \geq \widehat{\delta}$, where $\widehat{\delta} > \overline{\delta}$ denotes the discount factor above which monopoly pricing is the best achievable collusive price where no firm has an incentive to deviate. This relationship implies that collusion is easier to sustain the fewer the number of firms there are in the market. To derive a prediction on what happens to price-response asymmetry when the number of firms changes, I focus on what happens when $\delta = \overline{\delta}$. In the HH model, the only sustainable collusive price path when $\delta = \overline{\delta}$ is one where profits don't change as costs change. Under Bertrand competition, this would imply that markups must remain constant as costs change, or that the price-response asymmetry cannot exist when $\delta = \overline{\delta}$. The asymmetry only occurs when $\delta > \overline{\delta}$. Provided there is continuity, price-response asymmetry will be increasing in δ in the range $\left(\overline{\delta}, \widehat{\delta}\right)$. Since $\overline{\delta}$ is increasing in n, it follows that the range of asymmetry-inducing δ values is decreasing in n. In other words, we are more likely to observe price-response asymmetry when there are fewer firms in the market.

Although the relationship between price-response asymmetry and the number of firms is inuitively appealing, it doesn't actually provide an indication of what should happen to price-response asymmetry as market power changes. This is because firm-level market power does not exist in the RS and HH models as described, regardless of the number of firms in the market. For there to be market power under Bertrand-style competition, we must introduce product differentiation into the model.

A number of papers have augmented the basic supergame model of tacit collusion to introduce product differentiaion. In general, this literature concludes that greater degrees of product differentiation lower the critical discount factor above which collusion is sustainable.[6] Put another way, the basic finding of this particular literature is that collusion is

[6] This relationship is not always monotonic. Specifically, Ross (1992) derives a U-shaped relationship between product differentiation and the critical discount *rate* when inverse demand is linear in the quantity purchased of both goods. The critical discount rate equals one when the goods are perfect substitutes, and approaches one again when the goods have zero substitutability. For a large portion of Ross' measure of substitutability, the more homogeneous the goods, the less stable is the cartel. Ross (1992) also demonstrates that under a spatial, Hotelling-style model of preferences, the critical discount rate is monotonically decreasing in the degree of homogeneity, i.e., the more homogeneous the two products, the less stable is the cartel. Chang (1991) independently confirms that in a model of spatial competition, the critical discount

generally found to be easier to sustain the less homogeneous are the goods in the market.[7] Finally, to the extent that this result mirrors the earlier results regarding the effect of the number of firms on the critical discount factor and price-response asymmetry, we should expect that price-response asymmetry will decrease the more homogeneous are the firms in the market.

Empirically, the evidence in support of a market power effect on pricing asymmetry has been sparse. Borenstein, Cameron, and Gilbert (1997; henceforth BCG) attempt to empirically distinguish among several competing hypotheses for the source of the asymmetry. They find that the asymmetry between retail prices and terminal costs is consistent with a model of tacit cooperation where firms use the preceding period's price as a focal point. But the finding is not definitive, as their results are also consistent with a consumer search hypothesis. In the follow-up literature spawned by BCG, much of the attention is placed on identifying a market power effect at upstream levels in the gasoline industry, as in Borenstein and Shepard (2002), which finds that market power augments the asymmetry of wholesale cost responses to crude oil price changes.

An initial pass at the potential for a market power effect at a disaggregate retail level was attempted in Lewis (2003) by estimating an error correction model similar to the base model (Equation (4)) in this paper, but separately for the lowest and highest margin stations in the sample. Lewis finds no measurable market power effect. Deltas (2004) exploits markup differences across US states with a panel of monthly data to find evidence of a positive market power effect on asymmetry. The empirical approach in Deltas is directly analogous to the strategy in Lewis, although I argue later that it is not an approach I can follow here

factor is monotonically increasing as the products become closer substitutes. This result is shown to be robust to endogenous product location choice in Chang (1992), Hackner (1995), and Hackner (1996). Gupta and Venkatu (2002) argue that a common theme among the above listed papers is the assumption of mill pricing. They demonstrate that if firms instead use delivered pricing, then the result is reversed, i.e., the critical discount factor that sustains collusion is monotonically decreasing as firms locate closer together.

[7] This finding contrasts sharply with the conventional wisdom. Stigler (1964) and Carlton and Perloff (2005) both argue that product homogeneity should enhance the formation and stability of collusion, largely because it is presumptively easier to detect cheating in homogeneous product markets than in heterogeneous ones. The supergame literature (*supra* note 6) that explicitly incorporates product differentiation and finds that product heterogeneity aids collusive price setting implicitly assumes that firms can perfectly and costlessly detect cheating by rival firms.

with a station-level analysis.

3 Price-Response Asymmetry in Gasoline Markets

Estimation of a local market power effect on price-response asymmetry requires a disaggregate (preferably station-level), high frequency panel data set of retail prices, costs, and sales. I am aware of no publicly available data set that meets this criterion at the station-level, so from July 2002 to May 2003, I collected weekly price observations for the South Orange County region of California.[8] I augmented this data set with block-group data from the 2000 Census, which includes information on household incomes, housing values, commute statistics, and other relevant local demand- and cost-proxy variables. I also collected physical features of the stations such as lot size, number of pumps, the presence of a convenience store and other characteristics. New Image Marketing, Ltd. graciously provided me with station-level information on the operation types (Jobber, Lessee-Dealer) for a portion of the population, and I was able to complete these data by interviewing station managers for the remainder of the population. For this part of the study I use publicly available spot prices on Los Angeles reformulated gasoline as a measure of wholesale costs.[9] Lastly, I calculated geographic distances between stations by collecting latitude and longitude information with a GPS unit.

As shown in Figure 2, the geography of South Orange County makes it an ideal study area for this analysis. The region exhibits a nearly complete natural market boundary, which helps to avoid the potential arbitrariness of traditional market boundaries at, say, a specified street or arterial when there may be stations on the excluded side of the boundary that are valid rivals.[10] Additionally, according to the 2000 Census, approximately 99% of

[8] Lundberg Survey offers the closest match for this study's data requirements. Unfortunately, I found that Lundberg Survey is currently unwilling to make its data publicly available, even for sale.

[9] I was also able to obtain weekly rack prices by brand for the study period. Unfortunately, these data are difficult to use at the station level because Arco and Mobil, which together comprise approximately 1/3 of the market, do not participate in the rack market. In line with their organizational structures, Arco and Mobil sell directly to stations at Dealer Tank Wagon prices that are generally not available to the public (cf earlier footnote regarding Lundberg Survey).

[10] The study area is bordered by the Laguna Canyon wildlife preserve to the north, the Cleveland National Forest to the East, Camp Pendleton Marine Base to the south, and the Pacific Ocean to the West. While

the South Orange County residents work more than 5 miles from their homes, while over 40% commute more than 25 miles from their homes, with the predominant work destination lying in Central/North Orange County and in Los Angeles County. That the majority of commuters are traveling outside of the main study region for work has an important implication for competition between stations within the region. Competition between submarkets in the study area should be highly localized, in that it may be reasonable to assume that stations within the region are competing against each other primarily for those consumers who purchase gasoline near their homes. While there is a relevant outside good beyond the study region, consumers should prefer it relative to their own local stations in a roughly equal manner throughout the study region. As a result, the following analyses model prices as if I were dealing with the entire population of relevant stations, assuming away any location-specific biases that might arise if some stations were more price-sensitive to the outside good than others.

The study period itself covers a combination of cost shocks that resulted in a run-up of wholesale and retail prices in late 2002 and early 2003, as well as a sharp decline in wholesale costs in April of 2003. Average retail prices and wholesale spot prices are charted in Figure 3, which illustrates the December-March price run-up, the April-May decline and stabilization of prices, as well as a cyclical November decline in wholesale costs. Additionally, the charted prices demonstrate the underlying asymmetry, with retail prices appearing to rise much faster with cost increases than they fall with cost decreases. I exploit these data series, in addition to variation in retail prices at the station level, to estimate the asymmetric relationship between wholesale cost and retail prices changes. Table 1 provides summary statistics for these data, first at average levels, then broken down by positive and negative one-period changes.

---

there are gasoline tanks and fueling stations at Camp Pendleton, none are available to the general public.

3.1 An empirical model of price-response asymmetry

In order to model the dynamic nature of gasoline prices and their relationship to wholesale costs, I initially consider an autoregressive distributed lag (ARDL) model of the form:

$$p_{st} = \alpha_s + \sum_{j=0}^{J+1} \beta_{js} c_{t-j} + \sum_{k=1}^{K+1} \gamma_{ks} p_{s,t-k} + u_{st}, \tag{1}$$

where p_{st} denotes the retail gasoline price for station s at time t, c_{t-j} is the wholesale spot price for gasoline at period $t-j$, and u_{st} is assumed to be white noise.[11] By construction, (1) allows for short-term fluctuations in price levels as a function of recent prices and costs. Yet economic theory informs us that prices and costs should be governed by a long term relationship that sees prices increasing with costs. If we fix costs at some level $\overline{c}$, so that $c_t = c_{t-1} = ... = c_{t-J-1} = \overline{c}$, and take the conditional expectation of (1) , we get (asumming stationarity and ergodicity)

$$\begin{aligned} E\left(p_{st} | \{c_t\} = \overline{c}\right) &= \alpha_s + \overline{c} \sum_{j=0}^{J+1} \beta_{js} + \sum_{k=1}^{K+1} \gamma_{ks} E\left(p_{s,t-k} | \{c_t\} = \overline{c}\right) \\ &= \alpha_s + \overline{c} \sum_{j=0}^{J+1} \beta_{js} + E\left(p_{s,t} | \{c_t\} = \overline{c}\right) \sum_{k=1}^{K+1} \gamma_{ks}, \end{aligned}$$

or that

$$\begin{aligned} E\left(p_{st} | \{c_t\} = \overline{c}\right) &= \frac{\alpha_s}{1 - \sum_{k=1}^{K+1} \gamma_{ks}} + \frac{\sum_{j=0}^{J+1} \beta_{js}}{1 - \sum_{k=1}^{K+1} \gamma_{ks}} \overline{c} \\ &= \widetilde{\alpha}_s + \theta_s \overline{c}. \end{aligned} \tag{2}$$

This representation has several appealing features. First, it describes the long-run relationship between expected prices and costs. But moreover, it also gives the parameters of the model simple economic interpretations. If we define the markup between prices and costs as $m_{st} = p_{st} - c_t$, then the expected markup is $E\left(m_{st} | c_t = \overline{c}\right) = \widetilde{\alpha}_s + (\theta_s - 1)\overline{c}$. When combined with the derivative of (2) with respect to $\overline{c}$, $\partial E\left(p_{st} | \{c_t\} = \overline{c}\right) / \partial \overline{c} = \theta$, then,

[11] Duffy-Deno (1996) provides a wide summary of the empirical literature on price-response asymmetry to wholesale cost changes in the gasoline industry, and the different methods applied to estimate the relationship. At a high level, there seem to be two general approaches: partial adjustment and error correction models.

conditional on some value for $\bar{c}$ we can elicit direct a priori knowledge of the parameters $\widetilde{\alpha}_s$ and θ_s. θ_s is the long-run response of expected prices to a one-time shift in costs, while $\widetilde{\alpha}_s$ is backed out as the intercept that rationalizes our information on expected station-level markups; in the unique case that $\theta_s = 1$, $\widetilde{\alpha}_s$ is exactly equal to the long-run expected markup.

Because the parameters $\widetilde{\alpha}_s$ and θ_s provide such direct interpretation, most empirical studies of the gasoline industry choose to reparameterize (1) so that $\widetilde{\alpha}_s$ and θ_s are directly estimated in the model.[12] By working with the identities $\Delta p_{st} = p_{st} - p_{s,t-1}$ and $\Delta c_t = c_t - c_{t-1}$, we can judiciously substitute for p_{st} and c_t in (1) to derive the Error Correction Model (ECM) as:[13]

$$\Delta p_{st} = \sum_{j=0}^{J} \widetilde{\beta}_{sj} \Delta c_{t-j} + \sum_{k=1}^{K} \widetilde{\gamma}_{sk} \Delta p_{s,t-k} + \lambda_s \left(p_{s,t-1} - \widetilde{\alpha}_s - \theta_s c_{t-1}\right) + u_{st}, \tag{3}$$

where $\lambda_s = -\left(1 - \sum_{k=1}^{K+1} \gamma_{ks}\right)$, $\widetilde{\alpha}_s$ and θ_s are defined as above, and $\widetilde{\beta}_{sj}$ and $\widetilde{\gamma}_{sk}$ are specific to the number of initial price and costs lags in the levels specification.

After reparameterization, the remaining parameters in the ECM are also easily interpreted. The coefficients on the changes in spot prices, $\widetilde{\beta}_{sj}$, represent the short-run price adjustments in retail prices to changes in costs for station s, and the $\widetilde{\gamma}_{sk}$ coefficients represent short-run responses for station s to own-changes in retail prices during the last K periods. The error correction term, $p_{s,t-1} - \widetilde{\alpha}_s - \theta_s c_{t-1}$, can be thought of as the (one period lagged) deviation in prices from their long-run expected relationship with costs, which means that λ_s is the short-run correction in current prices that helps bring retail prices back to their equilibrium relationship with costs. When prices exceed costs by more than the

[12] See, for example, BCG, Balke et al. (1998), and Lewis (2003).

[13] For example, if J=K=2, then

$$p_{st} = \alpha_s + \beta_{0s} c_t + \beta_{1s} c_{t-1} + \beta_{2s} c_{t-2} + \beta_{3s} c_{t-3} + \gamma_1 p_{s,t-1} + \gamma_2 p_{s,t-2} + \gamma_3 p_{s,t-3} + u_{st}.$$

We then make the following substitutions: plug in $p_{s,t-1} + \Delta p_{st}$ for p_{st} on the left, and on the right substitute $c_{t-1} + \Delta c_t$ for c_t, $c_{t-1} - \Delta c_{t-1}$ for c_{t-2}, $c_{t-1} - \Delta c_{t-1} - \Delta c_{t-2}$ for c_{t-3}, $p_{s,t-1} - \Delta p_{s,t-1}$ for $p_{s,t-2}$, and $p_{s,t-1} - \Delta p_{s,t-1} - \Delta p_{s,t-2}$ for $p_{s,t-3}$. With these substitions, we get

$$\Delta p_{st} = \widetilde{\beta}_{0s} \Delta c_t + \widetilde{\beta}_{1s} \Delta c_{t-1} + \widetilde{\beta}_{2s} \Delta c_{t-2} + \widetilde{\gamma}_{1s} \Delta p_{s,t-1} + \widetilde{\gamma}_{2s} \Delta p_{s,t-2} + \lambda_s \left(p_{s,t-1} - \widetilde{\alpha}_s - \theta_s c_{t-1}\right) + u_{st},$$

where $\widetilde{\beta}_{0s} = \beta_{0s}$, $\widetilde{\beta}_{1s} = -\beta_{2s} - \beta_{3s}$, $\widetilde{\beta}_{2s} = -\beta_{3s}$, $\widetilde{\gamma}_{1s} = -\gamma_{2s} - \gamma_{3s}$, $\widetilde{\gamma}_{2s} = -\gamma_{3s}$, $\lambda_s = \gamma_{1s} + \gamma_{2s} + \gamma_{3s} - 1$, $\widetilde{\alpha}_s = -\alpha_s/\lambda_s$, and $\theta_s = -\left(\beta_{0s} + \beta_{1s} + \beta_{2s} + \beta_{3s}\right)/\lambda_s$.

long-run markup, we would expect downward pressure on current prices until the long-run markup is restored, and thus I expect λ_s to be negative.

Estimation of the ECM in traditional studies is usually performed in two stages: first, regress p_{st} on c_{t-1} and a set of station dummies, then second, construct $v_{s,t-1} = p_{s,t-1} - \widehat{\alpha}_s - \widehat{\theta}_s c_{t-1}$ and substitute these residuals in place of the error correction term in a least squares regression of (3).[14] Alternative practices are to perform unrestricted least squares (see BCG and Balke et al. 1998) of (3) after rewriting it to be linear in the covariates, or simultaneous one-stage estimation on (3) directly (see Lewis 2003) via nonlinear least squares or maximum likelihood. In this paper I advocate joint estimation of all the parameters in the context of a nonlinear Bayesian regression.[15] Earlier studies have relied on an assumption that if demand is linear in the long-run and marginal costs are truly constant across quantity, there should be 100% pass-through of costs to prices, implying that θ_s should be close to 1. Prior information on θ_s can range anywhere from relatively diffuse, allowing the data to reveal the implied long-term trend, to very dogmatic at $\theta_s = 1$ with certainty, analogous to the final approach taken in Lewis (2003) after the one-step procedure there revealed an estimated $\theta_s > 1$.

Asymmetry in the ECM is introduced by allowing separate coefficients for positive and negative changes in retail prices and wholesale costs. Letting $\mathbb{I}(\cdot)$ denote the indicator function, which takes on the value 1 when the interior condition $(\cdot)$ is met, a parsimonious

[14] This procedure is generally adopted out of concern that retail spot prices and wholesale costs are cointegrated, which occurs when the data generating process for wholesale costs is a unit root process and there is a 1:1 mapping between prices and costs. When true, this implies that prices are regressed on nonstationary covariates in the ARDL model, so that the standard sampling properties of the OLS estimator are no longer applicable. The two-step procedure in the ECM avoids this issue, since $v_{s,t-1}$ is stationary once we "know" $\widetilde{\alpha}_s$ and θ_s, and there are no longer any nonstationary regressors in the model. See Engle and Granger (1987) for a discussion of the asymptotic efficiency of this two-step procedure.

[15] There is an important identification problem for one-step and likelihood-based estimation of the ECM. At $\lambda = 0$, the parameters in the error correction term ($\widetilde{\alpha}_s$ and θ_s) are unidentified. One method for handling this would be to perform some form of restricted nonlinear least squares or restricted maximum likelihood estimation where λ is allowed to take on, say, only negative values. This would be analogous to a Bayesian procedure that placed zero prior mass on $\lambda \in [0, \infty)$. Bauwens et al. (1999) propose a joint prior between λ and $(\widetilde{\alpha}_s, \theta_s)$ that places zero prior probability on the event $\lambda = 0$. Another alternative is to allow the local non-identification and to just use a proper prior on λ, $\widetilde{\alpha}_s$, and θ_s; so long as the prior is proper, even if the likelihood does not identify the parameter, the posterior will still be a proper distribution. Moreover, provided the likelihood favors values away from the locally nonidentified point ($\lambda = 0$), learning will occur and the posterior will be updated.

representation of response asymmetry is

$$\begin{aligned}\Delta p_{st} = &\sum_{j=0}^{J} \Delta c_{t-j}\left(\widetilde{\beta}_{sj}^{-} + \widetilde{\beta}_{sj}^{+}\mathbb{I}\left(\Delta c_{t-j} > 0\right)\right) + \sum_{k=1}^{K} \Delta p_{s,t-k}\left(\widetilde{\gamma}_{sk}^{-} + \widetilde{\gamma}_{sk}^{+}\mathbb{I}\left(\Delta p_{s,t-k} > 0\right)\right) \\ &+\lambda_s\left(p_{s,t-1} - \widetilde{\alpha}_s - \theta_s c_{t-1}\right) + u_{st}.\end{aligned} \tag{4}$$

Under this construction, for a negative unit cost shock, the short-run contribution to the change in price is $\widetilde{\beta}_{sj}^{-}$, while for a positive cost shock, the contribution is $\widetilde{\beta}_{sj}^{-} + \widetilde{\beta}_{sj}^{+}$.

Past studies have illustrated the estimated asymmetry graphically via the difference in cumulative response functions (CRFs).[16] The CRF posits a single shock in wholesale costs at time t, and subsequently describes the path of cumulative price changes until prices settle at the equilibrium value associated with the new cost level. Without loss of generality, set $p_{st} = \widetilde{\alpha}_s$ and $c_t = 0$ (so that prices are currently at their long-run equilibrium), and then let $\Delta c_{t+1} = 1$. In addition, since we want to plot the CRF in the positive domain for both positive and negative shocks, assume that Δc_{t+1} is just the magnitude of the cost shock. Then the change in retail prices at $t+1$ for a negative shock is $\Delta p_{s,t+1} = \widetilde{\beta}_{s0}^{-}$. At $t+2$, the cumulative price change is $p_{s,t+1} + \widetilde{\beta}_{s1}^{-} + \Delta p_{s,t+1}\left(\widetilde{\gamma}_{s1}^{-} + \widetilde{\gamma}_{s1}^{+}\mathbb{I}\left(\Delta p_{s,t+1} > 0\right)\right) + \lambda_s\left(p_{s,t+1} - \theta_s\right).$ The pattern is relatively simple—at each period, the CRF is just the previous period's predicted price plus the predicted change in prices in the current period:

$$CRF_{s,t+f}^{-} = p_{s,t+f-1} + \widetilde{\beta}_{s,f-1}^{-} + \sum_{k=1}^{K} \Delta p_{s,t+f-k}\left(\widetilde{\gamma}_{s,k}^{-} + \widetilde{\gamma}_{s,k}^{+}\mathbb{I}\left(\Delta p_{s,t+f-k} > 0\right)\right) + \lambda\left(p_{s,t+f-1} - \theta\right). \tag{5}$$

An analogous CRF arises for a positive cost shock, with the only change that we add $\widetilde{\beta}_{s,f-1}^{+}$ to $CRF_{s,t+f}^{-}$. Lastly, I define an asymmetry function, $A_{s,t+f}$, as the difference between the positive and negative CRFs, i.e.,

$$A_{s,t+f} = CRF_{s,t+f}^{+} - CRF_{s,t+f}^{-}. \tag{6}$$

Note that this difference is not just $\beta_{s,f-1}^{+}$, since the price response for each cost change (positive or negative), at each period, depends on the past history of price changes as well.

In the sections that follow, I first establish the presence of asymmetry in the data across

[16] This is the approach taken in BCG and Lewis(2003).

stations. To implement this basic version of the model, I specify that all of the station-specific effects, i.e., $\widetilde{\beta}_{sj}$, $\widetilde{\gamma}_{sj}$, λ_s, and θ_s are constant across stations (e.g., $\theta_s = \theta \; \forall \; s$), and that station-level markups are fixed constants. After establishing this basic asymmetry, I then explore the data further by introducing random effects for each of the coefficients and interact them with the observed station-varying characteristics in my data. The random coefficients model is explained in more detail later.

3.2 Prices respond asymmetrically to cost shocks

For all of the coefficients in the model, I find that the posterior distributions are dominated by the likelihood, generally with high posterior precision. Summary output of the marginal posterior distributions is offered in Table A1, but not discussed here for brevity and ease of exposition. A discussion of the prior specification employed for this analysis is also offered in the appendix.

Pricing asymmetry is evidenced by the estimated CRFs in Figure 4a. We see that during the first 8 weeks after a wholesale cost shock, with the exception of week 2, retail prices respond more strongly to a cost increase than they do to a cost decrease. This finding appears to be statistically strong, i.e., precisely estimated, as well—the 95% equal-tail probability intervals for each CRF do not overlap save for the exception at week 2 and after week 8, where the two CRFs converge at the long run average effect. The posterior asymmetry function, illustrated in Figure 4b, also demonstrates this finding, with 95% of its mass centered on values greater than zero for all periods except week 2. The sharp change in the second week following a cost shock indicates that, on average, stations in my sample temporarily but significantly slow down the rate of response to cost increases during the second week, while for a cost decrease, the temporary flattening of the response rate occurs during the third week.

Regarding estimation of the long-term relationship between retail prices and wholesale costs, I find that the system at first appears potentially nonstationary. The coefficient on lagged costs in the error correction term is centered at 1.38 with a standard deviation of 0.02

and a posterior probability that it exceeds unity arbitrarily close to 1. This occurs despite the nearly dogmatic prior imposed on this coefficient, which is centered at 1 with a standard deviation of 0.1/6. One explanation of this result is that I work with a relatively short sample period relative to other studies, the end of which sees average markups settling at a value greater than they were initially at the start of the sample. Yet, while BCG estimate this coefficient to be approximately 1 with relatively small standard error, Lewis (2003) also finds a value larger than 1 under one-stage estimation ($\widehat{\eta} \approx 1.5$ with standard error 0.29 in that study), despite having 92 weekly observations compared to my 39. On balance, it would appear that the data strongly favor $\theta > 1$, especially in light of the of the data significantly updating my highly informative prior specification.

One response to this apparently strong result is to follow Lewis' study and reestimate the model under a still more dogmatic prior, perhaps by setting $\theta = 1$ with certainty. Instead, I think the appropriate response in this study is to accept the data as properly informative about the nature of the residual demand facing each station. A possible alternative economic explanation of $\theta > 1$ to the statistical one just described is that each station may exhibit constant elasticity for it's residual demand (as opposed to linear residual demand). Bulow and Pfleiderer (1983) show that under constant elasticity of demand, the price response to a change in marginal cost is $\partial p/\partial c = \eta/(\eta + 1)$, where for a firm that faces imperfect substitutes the elasticity η is less than -1. While I would stress that this is a highly stylized assumption, if indeed the stations in my dataset face constant elasticty residual demand, the implied average station-level elasticity when $\theta = 1.38$ is $\eta = -3.63$. This number is consistent with the elasticities found in Barron et al. (2003), but considerably less elastic than the findings in Romley (2002).

4 Spatial Differentiation and Price-Response Asymmetry

The most direct approach to exploring a local market power effect on price-response asymmetry is the procedure adopted by Deltas (2004), which interacts observed markups with the Δc_{t-j} and Δp_{t-k} terms in the ECM. While this strategy is efficacious at the aggregate level, it can lead to spurious conclusions at the station level. At an aggregate level, the unit of observation is the average retail price for a given region. To the extent that the majority of regions studied do not differ drastically in their organizational makeup, the average cost across regions will vary across regions largely as a function *only* of regional-specific wholesale spot prices, which are readily observable to the researcher.

But at the station-level, the observed markup of retail price minus wholesale cost is not the true markup. The actual marginal cost faced by a station varies across stations. In general, I expect a station's marginal cost to depend proportionally on the wholesale spot price of gasoline, i.e.,

$$c_s^* = a_s + b_s c.$$

While the spot price of gasoline is constant for all stations in a common regional market, the marginal costs across stations will not be. For example, suppose, as this study finds, that salary-operated (company owned and operated) stations have a markup over spot prices that is less than that for lessee-dealer stations. One would be tempted to conclude, by the Lerner index measure of market power $l = (p - c) / p$, that salary-operated stations exhibit lower local market power than lessee-dealer stations. This conclusion is premature—if the industry wisdom is correct that salary-operated stations have lower marginal costs than lessee-dealer, then, although the salary-operated stations charge lower prices, their markups may not be lower than those of lessee-dealer stations, and could even be greater.

Ideally, if I had cost data that was a better approximation of marginal costs at the station level—say, rack (terminal) price data for the independent stations *and* Dealer Tank Wagon (DTW) price data for the company owned stations—then I could follow the procedure employed by Deltas (2004) and compare the implied pricing asymmetries across stations with observed differences in markups.[17] Unfortunately, DTW price data proved to be unavailable for public sale. With 40% of my stations characterized as lessee-dealer, if I were to instead proxy station-level marginal costs with rack price data, I would be introducing a

[17] DTW is the price a dealer charges a station for both the gasoline and the delivery together. Company-owned, lessee-dealer stations purchase their gasoline at the DTW price.

non-ignorable sampling bias into my estimation of each station's markup.

Rather than introducing markups directly into the ECM as just suggested, I build on the evidence suggested in other studies of the gasoline industry that spatial differentiation can influence a station's local market power. By interacting station-level characteristics with the cost- and price-change variables in the ECM, I can then predict the influence of that characteristic on a station's price-response asymmetry. By association, I can then derive an implied relationship between market power and price-response asymmetry. For example, Romley (2002) finds that branding (specifically, upgrading to a Chevron station) can decrease a station's own-price elasticity from 11.4 to 8.8, which corresponds to a 3 percentage point increase in the Lerner index and therefore an increase in market power. A finding in this study that branded stations demonstrate greater price-response asymmetry than unbranded stations would thus suggest that not only do stations pass through cost increases faster than decreases, but local market power achieved through brand differentiation allows this difference in pass through rates to be greater for branded stations than it is for unbranded.

I explore the influence of spatial differentiation on price-response asymmetry by considering a wide set of station-level characteristics. Section 3 describes the data set and its collection procedure. From these data I selected a subset of characteristics that I think are most likely to influence a station's market power. Among those selected are brand identity, the presence of a carwash, service station, or convenience store, and the station's lot size. Additionally I look at demand shifters that can be proxied by market demographics such as local household income and the size of the local population.[18] Other characteristics that I consider are site-specific features such as distance from the nearest major freeway, the density of pumps (pumps per acre), and ease-of-access variables such as the number of driveways, whether any of the driveways has a traffic light, and whether the primary arterial is divided or not.[19] I also look at other local market characteristics such as whether

[18] I determine the values for local market demographics by associating all of the block groups within 1 mile of a station as comprising the relevant local market. I then construct the average value of the characteristic across the selected block groups.

[19] Much of South Orange County is newer, planned development, and a common street design is to include

the land use in immediate proximity to the station is residential or commercial, whether a competitor's pricing is directly visible from the station, and whether the station itself is contained in a shopping center. Lastly, I include the station's organizational relationship with its parent refinery as a proxy for a cost shifter.[20] Summary statistics and the incidence of these site and local market characteristics are described in Table 2.

4.1 Empirically estimating station-level variation in price-response asymmetry

I argue above that by interacting spatial characteristics with the short- and long-run changes in costs and prices, I can predict separate asymmetric relationships associated for each characteristic while controlling for the effect of other spatial features. I accomplish this in the ECM model specified in equation (4) by allowing each station to have its own random coefficient on the cost and price change variables. I then describe the station-level variation in the coefficients with the following set of assumptions:

$$\begin{aligned}\widetilde{\beta}_s &= \left(\widetilde{\beta}^-_{0,s}, \widetilde{\beta}^+_{0,s} ... \widetilde{\beta}^-_{J,s}, \widetilde{\beta}^+_{J,s}, \widetilde{\gamma}^-_{1,s}, \widetilde{\gamma}^+_{1,s}, ... \widetilde{\gamma}^-_{K,s}, \widetilde{\gamma}^+_{K,s}\right)' \\ &\overset{ind}{\sim} \mathcal{N}\left(\left(I_{J+K} \otimes w_s\right)\beta^*, \Sigma_\beta\right)\end{aligned}$$

$$\lambda_s \overset{ind}{\sim} \mathcal{N}\left(w_s \lambda^*, \sigma^2_\lambda\right)$$

$$\theta_s \overset{ind}{\sim} \mathcal{N}\left(w_s \theta^*, \sigma^2_\theta\right)$$

$$\widetilde{\alpha} = \left(\widetilde{\alpha}_1, ..., \widetilde{\alpha}_S\right)' \sim \mathcal{N}\left(W\alpha^*, A\right),$$

where, for notational convenience, I stack both the cost- and price-change coefficients into a single vector $\widetilde{\beta}_s$, associated now with the covariate vector $x_{st} = \left(\Delta c^-_t, ..., \Delta p^+_{s,t-K}\right)$. Under these assumptions, the model is now described as a hierarchical error correction model (H-ECM) that centers each covariate in (4) at a linear combination of the L observed station-level characteristics in the vector w_s. Hence, β^* is the $(J+K)\,L$ vector of coefficients in the *second-stage* analysis. Alternatively, if we integrate over the distributional assumption

highway dividers which restrict left-turn access into commercial centers.

[20]Generally speaking, stations can be independently owned and operated, jobber owned and operated, or company (refinery) owned and either lessee-dealer or company operated. Industry wisdom holds that company owned and operated stations have the lowest wholesale costs while independent and lessee-dealer stations have the highest wholesale costs.

on $\widetilde{\beta}_s$, β^* is the coefficient vector on the interaction terms between the covariates in x_{st} and w_s. Similarly, λ^*, θ^*, and α^* are the L vectors of coefficients for the remaining parameters.

The hierarchical setting of this model is a flexible and convenient method of introducing station-level variation in the first-stage regression coefficients. We can think of the hierarchical specification as implying a sort of two-stage regression, where we might first estimate a separate ECM equation for each station, and then subsequently regress each of the coefficients from the first stage on the corresponding station-level covariates. In contrast to this approach, the H-ECM estimates all of the parameters in the implied two-stage approach jointly. Additionally, it allows for *spatial* correlation to be specified in the station-specific effects. While the normality assumption may at first appear restrictive, it is also appealing in that it implies the marginal distribution of Δp_{st} is also normal when the residual u_{st} is normal, and is thus no more restrictive than the classical normal linear regression model with interaction terms and a specific form of heteroskedasticity.

In my data set, I observe a limited number of demand and cost variables at the station level. To the extent that differencing prices over time removes any time-invariant unobserved correlation across stations as a result of spatial differentiation, the ECM is appropriately modeled with independent covariates for the short-run cost and price changes. However, the long-run markups depend explicitly on each station's degree of spatial differentiation. Out of concern that the observed spatial characteristics in the data may not completely describe each station's long-run markup, in particular with regard to my local market demand proxies from the census data, I allow for spatial correlation in the markups as a function of the distance between stations. In the hierarchical specification above, this spatial dependence occurs in the covariance matrix A in the distribution of the station-specific markups $\widetilde{\alpha}_s$. Following Banerjee, Carlin and Gelfand (2004), I describe A with a simple exponential covariance function, where

$$A_{ij} = \begin{cases} \sigma_\alpha^2 \exp\left(-\phi_\alpha d_{ij}\right) & \text{if } i \neq j \\ \tau_\alpha^2 + \sigma_\alpha^2 & \text{if } i = j \end{cases} .$$

This specification for the covariance between stations implies that the spatial dependence is decreasing with the distance d_{ij} between stations, and is always nonnegative.[21]

4.2 Price-response asymmetry varies with a station's spatial characteristics

As mentioned above, my empirical strategy for identifying a local market power effect on pricing asymmetry is to allow station-specific price responses to a cost shock and then to compare the resulting predictive asymmetry functions across station-level characteristics. A summary table of the estimated effects of the covariates on retail gasoline markups is provided in Table A2, while Table A3 provides information on the posterior distribution of the short-run coefficient parameters. Unfortunately, the information contained in these two tables is somewhat overwhelming, and as a result I concentrate my discussion of the effect of the covariates on pricing asymmetry to the graphical analyses below.[22,23,24]

Still, there are some parameters in the model that are directly interpretable regarding the effect of station characteristics on pricing dynamics. In Table 3, I present a summary of the posterior distribution for λ^*, the coefficients in the "auxilliary regression" of the station-level λ parameters on the characteristics described in Table 2. To interpret this summary output, recall that we expect λ to be negative if stations exhibit a tendency to return to a long-run linear relationship between retail and spot prices, so that negative values for the coefficients λ^* indicate characteristics which speed up the return to equilibrium while positive values slow down this response. Among the results in Table 3 that indicate a market-power effect on station-level pricing dynamics are the effects of nearby competition and the presence of carwashes and convenience stores at a site. In general, it appears that most branded stations return to their long-run equilibrium price-cost relationships faster

[21] The nonnegative property of the exponential covariance function implies that the spatial correlation between stations is always positive. This is a significant limitation of most simple models of spatial dependence. While a more complicated specification may allow for a nonmonotonic covariance function in d_{ij}, as a practical matter the markups in my data set appear to be highly positively correlated.

[22] I consider an ECM model with $J = 3$ cost-change lags and $K = 2$ price-change lags, each of which has both a positive and a negative coefficient. Each of these 10 coefficients, plus the coefficients in the long-run relationship is associated with the 31 station-level characteristics described in Table 2, which results in 403 total coefficient parameters in the model.

[23] As with the basic model, I have also included details on the prior specification and estimation procedure in the appendix.

[24] See the discussion that follows for a possible explanation for the Arco and Mobil difference.

than unbranded stations, with the exception of Arco and Mobil.

Although a tabular approach has limited appeal in this paper, a diagrammatic one does illustrate well the predictive effect of the station characteristics on price-response asymmetry. In Figure 5, I provide the results of a predictive analysis for the effect of being a salary-operated (company-operated) station relative to being lessee-dealer operated. A nonlinear predictive analysis like that contained in the CRFs and asymmetry functions must condition on values for the other covariates yet also try to isolate the marginal effect of a particular characteristic. In order to achieve this, I determine the predictive CRF for all of the stations in my sample, and then generate a separate set of predictive CRFs for all of the stations but change the status of the salary-operated stations to lessee dealer, which gives me an as-if predictive result: if, all else being equal, the salary-operated stations were instead lessee-dealer, what would their cumulative response function look like? The average (across salary-operated stations) difference between the two CRFs and their associated asymmetry functions yields the marginal effect on the pricing dynamic of being a salary-operated station.

For example, in Figure 5a, I plot the mean of the predictive CRF for both a positive and negative cost shock for the salary-operated stations. An analogous plot is offered in panel b after changing their status to lessee-dealer. Panel c compares the posterior mean of the asymmetry functions for each type (salary-operated versus lessee-dealer), and panel d gives the full posterior distribution of the difference between the asymmetry functions. While panels a-c offer the most interesting picture of what is happening with prices across these two predictives, panel d yields the most important information on the difference in price responses between them. At each week following a single $1 cost shock I draw a boxplot for the posterior distribution of the difference between the predictive price-response asymmetry of salary-operated stations and their predictive asymmetry after switching them to lessee-dealer stations. The central point of the boxplot corresponds to the posterior median of this difference, which I have linked across weeks in the solid line connecting each boxplot. The upper and lower lines of the rectangular boxes in each boxplot correspond to the 75th and 25th percentiles of the distribution respectively, while the lines leading out of the box give an idea of the basic range of the distribution (note, they do not imply a 95% interval, which is entirely contained within this range).

The series of boxplots in Figure 5d therefore suggest that the point estimate of the difference in price-response asymmetry is positive in all weeks except week 2.[25] The negative difference in week 2 is also revealed in panel c, where we see the two mean predictive asymmetries crossing over at 2 weeks after the cost shock. I am also able to derive the probability that the difference in asymmetries is positive by looking at the amount of mass in each boxplot for each week that lies above 0. In weeks 1, 3, 4, 5, and 6, it appears that roughly 60% or more of the mass lies above zero, suggesting that the posterior probability that salary-operated stations have a wider asymmetry in price response to a cost shock than if they were lessee-dealer stations is at least 60% in those weeks.[26] While I would hesitate to describe this as definitive evidence that companies achieve greater price-response asymmetries with their salary-operated stations—for which they have complete control over the pricing decisions—than their lessee-dealer stations—for which they have only imperfect pricing control through DTW prices—it is suggestive of at least a small effect.

Yet it is not surprising that I am unable to precisely determine a marginal effect from being a salary-operated station, insofar as this is not a spatial dimension for which I expect customers to be concerned about or even aware of. I have included operation types in the analysis to control for cost variation in the data across stations, and present the results just described to illustrate the diagrammatic approach that I take in this paper to describe the predictive marginal effects. For the majority of the remaining station characteristics, I generally find highly suggestive evidence, and occasionally definitive evidence, that spatial differentiation does influence a station's price-response asymmetry.

[25] Under quadratic loss, the optimal Bayesian point estimate is the posterior mean, while under symmetric linear (absolute) loss, the optimal Bayesian point estimate is the posterior median.

[26] This is a subjective probability statement. It is the conditional probability that results from applying Bayes rule to update the prior distributions described in the appendix with the likelihood function described by the data and the hierarchical ECM model presented in this paper. In general (except for θ, which is discussed in Section 3), my prior distribution is highly diffuse and should be widely acceptable as a "public" prior.

The first major characteristic that I look at is the effect of brand identity on asymmetry. Figure 6a charts the predictive price-response asymmetry for 4 groups of brand identities in the data. I separate Arco and Mobil from the other branded and unbranded stations because they each have unique operating structures that distinguish them from the other brands. Specifically, both tend to prefer Lessee-Dealer or salary-operated contracts with their stations, largely avoiding the jobber and independent organizational types. In the case of Mobil, it appears that the resulting greater control over pricing leads to an early spike in the price-response asymmetry, with retail prices rising much faster in the first week after a positive cost shock than they fall after a negative cost shock. As with the overall average asymmetry, the asymmetry vanishes in the second week, but reappears a week later and diminishes toward zero afterward.

In Figure 6b, I call attention to the difference in price-response asymmetries between the branded (Chevron, Shell, and Unocal 76) and the unbranded stations. According to the predictive densities, it appears that branded stations have a far greater price-response asymmetry than the unbranded stations in the first 9 weeks after a cost shock, after which point the unbranded stations have a greater, albeit small in magnitude, asymmetry than the branded stations. In monetary terms, for branded stations the difference peaks three weeks after a cost shock, where prices after a cost increase tend to rise by more than 30 cents greater than they fall for a corresponding negative shock For unbranded stations, this difference is less than 15 cents, and the difference between the two groups is estimated at nearly 20 cents.

If we interpret the asymmetry as a cost the consumer bears by frequenting a particular station, then the results suggest with almost 90% certainty that consumers pay more relative to a cost decrease in the first several weeks after a cost shock by purchasing from branded stations than they would pay if they purchased from unbranded stations. Moreover, if the underlying process by which the asymmetry occurs is via implicit collusion, it would appear that it breaks down much more rapidly for unbranded stations than branded, consistent with the notion that tacit collusion is easier to maintain for stations with more relative market power.

In Figure 7, I look at the issue of spatial differentiation more directly by looking at the benefit of geographic isolation for a station. Specifically, I identify those stations which have no rival stations within 0.1 miles, which in this data set amounts to those stations which do not share an intersection with another competitor. The solid line in Figure 7a plots the mean of the predicted asymmetry from a \$1 cost shock for these stations. I then consider the predicted outcome on this asymmetry from adding a rival station within 0.1 miles, which I plot as the dotted line Figure 7a. For the first two weeks after a cost shock, both geographic types exhibit a similar predicted asymmetry. But after the third week, the asymmetry for the station with an additional immediate rival is lower by approximately 7 cents. Even more revealing is the relative precision with which this difference in asymmetry is estimated. In Figure 7b, we see that after week 3 the majority of the posterior predictive mass lies above zero. In fact, the posterior probability that the effect of immediate isolation for these stations is positive exceeds 90% in weeks 4 through 12. As with the positive effect of branding on price-response asymmetry, the results on the benefits of isolation also suggest that market power increases the tacitly collusive equilibrium price if the underlying mechanism driving the asymmetry is tacit cooperation among station managers.

I also look at the marginal effects on price-response asymmetry of site-specific characteristics in Figure 8. Perhaps surprisingly, the results suggest that while there is a predictive positive difference in asymmetry for a station located one standard deviation (about 2.5 miles in my data) further from the nearest open-access freeway for all post-shock weeks after the first, the difference is never significant in the sense of statistical precision except perhaps for 2, 3, 5, 6, and 7 weeks after, and never strongly significant in the economic sense. Having a traffic light for at least one driveway appears to increase the predictive asymmetry, although the difference in the first 2 weeks after a cost shock probably offsets the high difference in the following couple of weeks from a cost-to-the-consumer perspective. Still, it appears that easier access into and out of a station is a notable dimension of spatial differentiation, with higher degrees of asymmetry of at least a few cents and as many as

15-20 cents persisting after the third week following a cost shock.

Stations which bundle a convenience store with their gasoline business also appear to spatially distinguish themselves from their competitors, although as with a driveway traffic light, it appears that there is an offsetting effect in the first few weeks, after which point the higher asymmetry associated with convenience store stations persists and is significant in the statistical sense. Together with the result for driveway traffic lights, it appears that stations with these characteristics are better able to maintain higher prices after a cost decrease than stations which do not have these characteristics, suggesting that collusive pricing may be less stable for the set of stations without these characteristics.

A stronger economic finding occurs with stations that have a one standard deviation higher density of pumps on their lot than other stations. To the extent that this is indicative of consumers' higher preferences for shorter wait times, this result again suggests a positive impact of spatial differentiation on price-response asymmetry. But it should be noted that estimating the coefficients on the number of pumps (per acre) is complicated by an endogeneity problem of the usual supply and demand kind. If stations enjoy any kind of volume discount, then their per-unit cost may be falling in the volume of customers they can serve, which is probably directly related to the number of pumps they have at their site. Still, regardless of the appropriate dimension, the finding is economically and statistically strong that price-response asymmetry is increasing in relative pump density.

Figure 9 describes the results of the predictive analysis for some of the local-market characteristics. Panel a offers an especially interesting comparison with respect to the tacit collusion theory described in Section 2. When rival stations are close enough to have their prices visible from each other's stations, it appears that higher collusive pricing may be easier to maintain in the first couple of weeks following a cost shock, but that afterward the benefit of price visibility in maintaining higher collusive prices disappears. Somewhat surprisingly, I also find a negative effect on asymmetry for being located in a shopping center, where one might suppose there would be higher consumer demand. A possible explanation for this is that when a station is located in a shopping center, there tends to be another one in close proximity, even if it is not directly visible.

I also include local demographics in Figure 9. The result in panel c that stations located near block groups with larger population sizes tend to have a wider price-response asymmetry is expected conditional on the hypothesis that spatial differentiation increases a station's relative market power. The peculiarity is in panel d with the predicted marginal effect of an increase in local household income. All else held equal, I would expect that if greater local market power does increase a station's price-response asymmetry, then being located to consumers with more income should widen the asymmetry, not shrink it.

A possible economic explanation for a negative income effect derives from the findings in previous studies, particularly Barron et al. (2003), that higher-grade gasoline is more price elastic than lower grade, and that substitution effects dominate over income effects with regard to gasoline pricing. The idea is that high-income consumers typically are more likely to buy high grade gasoline, but when prices rise sharply switch from high grade to low grade gasoline. This finding may extend to the broader gasoline market, so that higher income consumers are more likely in general to shop around for lower prices when prices are high, and thereby making it more difficult to maintain tacitly collusive pricing when wholesale prices fall.

On balance, the predictive results suggest that those local market and site characteristics which increase a station's local market power also tend to widen its price-response asymmetry. Effective station characteristics, in particular brand identity and larger local population sizes, and to a lesser extent improved ease-of-access and the offering of a convenience store, are associated with faster cost pass through when costs increase and slower cost pass-through when costs decrease. The finding that having a competitor in close proximity also shrinks the price-response asymmetry adds further credence to the suggestion that the mechanism by which local market power influences asymmetry may be related to the concept that markets composed of spatially close stations form less stable collusive regimes than markets composed of spatially distant stations.

5 Concluding Remarks and Directions for Future Research

In this study I examine the potential influence of spatial differentiation, and by extension, local market power, on the well documented empirical phenomenon that gasoline prices rise faster for a cost increase than they fall for a comparable cost decrease. Using a highly detailed station-level data set, I establish in this paper that price-response asymmetry is a dominant feature of the data. I then show that stations with specific site and local-market characteristics are associated with higher price-response asymmetry than stations without (or with lower levels of) these characteristics. To the extent that these spatial characteristics increase each station's potential local market power, the results suggest that market power does indeed augment price-response asymmetry.

These results also indicate a possible direction for future analyses of price-response asymmetry. A direct approach to measuring the effect of market power explicitly on price-response asymmetry would be to compare the asymmetry across different observed levels of market power. One of the most common measures of market power is the Lerner index, the creation of which without a structural demand model requires highly accurate cost data. This study circumvents the problem of inadequate station-level cost data by looking at spatial differentiation, which is assumed to positively influence a station's market power. Another option, which was attempted here but proved unsuccessful, is to acquire both DTW and rack pricing and construct a more accurate estimate of each station's cost. Alternatively, the econometric literature on the errors-in-variables problem with spot or rack prices as an inaccurate but positively correlated proxy for marginal costs could yield useful results. To have any statistical precision, such an approach would require a set of strong instruments to identify the correct cost effects and, by extension, the implied markups.

A further alternative is to combine the procedure here with a structural demand model that backs out own-price elasticities for each station, which would provide another method for estimating the Lerner index measure of market power. However, due to the common lack of good, or even any, quantity data at the station level, a natural inclination for the researcher attempting a structural analysis is to assert a static equilibrium model that identifies the underlying structural equations in spite of the missing quantity information. The complication facing this approach is that the industry has an obvious dynamic component. Indeed, the dynamic pricing of gasoline stations is the subject of study, and so the usual methods that achieve identification through a static game would seem inappropriate. While the static game might change over time in a systematic manner, it would be difficult to argue that the *static* game in period t was a function of outcomes or states in earlier periods. In this case, the appropriate structure is not a static model, but a dynamic one.

Appendix

A.1 Specifying a prior distribution for the basic ECM

For this stage of the analysis, there is a surprising wealth of prior information on the parameters of the model and the asymmetry I expect to see in the data a priori. Largely, this information derives from the existing literature, in particular BCG and Lewis (2003). The long-run parameters are the simplest to elicit. By way of construction, estimation of (the nonrandom parameters) $\widetilde{\alpha}_s$ and θ is eased by formation of a matrix composed of a constant, $S-1$ station dummies, and the vector of one-period lagged costs for each station. Thus $\widetilde{\alpha}_0$ is the markup for the excluded station, which I describe with a normal density, centered at 0.8 with a variance of 5; the remaining $S-1$ station dummy coefficients are centered at 0, also with a variance of 5. I interpret this prior information as being highly diffuse, centered at reasonable a priori values. Diffuseness, however, is an inappropriate specification for the long-run price response to costs represented in θ. BCG estimate θ at nearly 1 with small standard errors, while Lewis ultimately rejects the estimated values of θ and imposes $\theta = 1$ with certainty. I take a middle-of-the-road approach and continue to estimate θ, but with a highly informative prior centered at 1 and a variance suggested by the "six-sigma" rule, i.e., with a prior standard deviation of 0.1/6. In words, this translates to a prior belief that I am nearly certain (probability 0.99) that θ lies between 0.9 and 1.1.

For the short-run response parameters $\widetilde{\beta}$ and $\widetilde{\gamma}$, I revert to diffuseness with cues taken from BCG. I condition the analysis on $J = 3$ and $K = 2$, which results in 3 cost-difference regressors and 2 price-difference regressors. My prior on $\widetilde{\beta}_0$, $\widetilde{\beta}_1$ and $\widetilde{\gamma}_1$ is such that negative changes are centered at 0, while positive changes are centered at 0.1, reflecting an a priori expectation of a positive asymmetry difference. Uncertainty about the appropriateness of the additional lags leads me to center all of the coefficients in $\widetilde{\beta}_2$ and $\widetilde{\gamma}_2$ at zero. For all of the short-run changes in cost and price coefficients, the variance is set at 5, which again represents a diffuse specification that highly favors the likelihood information over the prior. For the short-run response to the deviation from equilibrium, λ, I specify a normal prior density centered at -0.5 with a variance of 1, suggesting a nearly certain prior belief that λ is negative and favoring a relatively quick return to the long-run equilibrium relationship between retail prices and wholesale costs.

Lastly, I assume that the residual in the regression, u_{st}, is $N\left(0, \sigma^2\right)$, and thus require a prior specification for the residual variance term in addition to the coefficients described above. Recall the usual specification for goodness of fit, $R^2 = 1 - SSE/SST$, where SSE is the sum of squared errors and SST is the total sum of squares. Also recall the classical estimator for σ^2, $s^2 = SSE/\left(N-L\right)$, where L is the total number of regressors. Combining this information one can rewrite $s^2 = \left(1-R^2\right) SST/\left(N-L\right)$. Thus given prior beliefs about goodness of fit and the underlying variance in the data, I center my distribution for σ^2 at my prior expectation of s^2. Additionally, since I specify an inverse gamma distribution for the residual variance, i.e., $\sigma^2 \sim IG\left(\nu/2, 2/\nu s^2\right)$, the minimum integer value that ν can be to ensure a prior mean is 3, which is the value I use in this study. Specifically, I center my prior information on the residual variance near the implied value for s^2 that results from $R^2 = 0.25$ and SST the actual variance in the data times the sample size. But since $\nu = 3$, this results in a prior distribution that, like those for the coefficients above, is highly diffuse. Indeed, it places so much weight in the tails of the distribution that the prior variance of σ^2 is not finite.[27]

A.2 Posterior estimation and output of the basic ECM

I estimate the ECM with asymmetry described in Section 3.1 via the Gibbs sampler, which allows me to obtain draws from the joint distribution of all the parameters using straightforward Bayesian linear regression techniques. This procedure has the added benefit that I estimate the long- and short-run parameters jointly without having to rely on the two-stage

[27] A valid critique of this prior elicitation procedure for σ^2 is that it is "data informed." Specifically, I condition on $SST = \Sigma_{st}\left(p_{st} - \overline{p}\right)$, which clearly depends on the price data. But the elicitation itself is very intuitive. It amounts to a prior expectation for both R^2 and the underlying variance in the data. The difficulty comes in that the data variance is a basic descriptive statistic that the researcher will always know before estimation of the model, and which is information he cannot forget when constructing prior beliefs about the model's fit. In a practical sense, this procedure results in centering the prior for σ^2 at no more than the underlying variance in the data—how much less than the underlying variance is presumably based on non-data information about R^2.

Engle and Granger (1987) procedure. As an added benefit, I also avoid having to rely on asymptotic results for the predictive analysis inherent in deriving the cumulative response and asymmetry functions—the posterior distribution of the parameters, and subsequently the posterior distributions of both CRF_{t+f} and A_{t+f}, is an exact, finite sample distribution. The Gibbs routine itself proceeds in the following manner:

1. Initialize $\widetilde{\beta}^{(0)}$, $\widetilde{\gamma}^{(0)}$, $\widetilde{\alpha}^{(0)}$, $\theta^{(0)}, \lambda^{(0)}$ and $\sigma^{2\ (0)}$. Then, for $m = 1, ..., M$:
2. Draw $\widetilde{\beta}^{(m)}, \widetilde{\gamma}^{(m)} | y,\ \widetilde{\alpha}^{(m-1)}, \lambda^{(m-1)},\ \theta^{(m-1)},\ \sigma^{2\ (m-1)}$
3. Draw $\lambda^{(m)}, \theta^{(m)} | y,\ \widetilde{\beta}^{(m)}, \widetilde{\gamma}^{(m)},\ \sigma^{2\ (m-1)}$
4. Draw $\widetilde{\alpha}^{(m)} | y,\ \widetilde{\beta}^{(m)}, \widetilde{\gamma}^{(m)}, \lambda^{(m)}, \theta^{(m)},\ \sigma^{2\ (m-1)}$
5. Draw $\sigma^{2\ (m)} | y,\ \widetilde{\beta}^{(m)}, \widetilde{\gamma}^{(m)}, \widetilde{\alpha}^{(m)}, \theta^{(m)}$

The simplicity of the Gibbs algorithm is particularly useful in this exercise, since conditional on the long-run parameters, $\widetilde{\beta}$ and $\widetilde{\gamma}$ are just the coefficients in a linear regression with known variance.[28] Likewise, conditional on the short-run parameters, $\widetilde{\alpha}$ and θ are also just coefficients in a linear regression with known variance. Hence each step in the Gibbs routine is no more difficult than standard Bayesian linear regression, and after convergence, all of the draws represent valid draws from the joint posterior distribution of the parameters. The results of the Gibbs routine for the basic asymmetry model that uses all of the stations are summarized below in Table A1.

A.3 Specifying a prior distribution for the hierarchical ECM

For the hierarchical ECM, I generalize the prior distribution described above to imply the same basic information about the first-stage coefficients as when they were nonrandom and equal across stations. For ease of exposition, let ξ_s denote an arbitrary coefficient in the first-stage regression (e.g., $\widetilde{\beta}_{0,s}^{-}$ or $\widetilde{\gamma}_{2,s}^{+}$). Recall the distributional assumption on these coefficients that the marginal density for one of the coefficients is $\xi_s \sim \mathcal{N}\left(w_s \xi^*, \sigma_\xi^2\right)$. In

[28] See Poirier (1995) for an explanation of estimation in the standard Bayesian linear regression model.

the original setup with $\xi_s = \xi$, when ξ was $\widetilde{\beta}_0^-$, $\widetilde{\beta}_1^-$ or $\widetilde{\gamma}_1^-$, I centered the prior at zero, while for the associated positive effects the prior was centered at 0.1. I maintain that specification through the prior on ξ^*. Since w_s contains a constant plus a set of station-level covariates, about which I want the data to be the primary source of information, I center ξ^* so that for the negative first-stage effects, $\xi^* = 0$, while for the positive first-stage effects, the first element of ξ^* (the coefficient on the constant in w_s) is 0.1, while the remaining elements are centered at 0. As with the basic model, I also specify a prior variance for ξ^* of 5, with no a priori expectation of covariance between each coefficient. Under a normal prior then, the vector of stacked coefficients from section 4.1 is $\beta^* \sim \mathcal{N}\left(\underline{\beta}^*, 5 * I_{L(J+K)}\right)$, where $\underline{\beta}^*$ has elements as just described above and I_n is the identity matrix of size n. Additionally, since I do not want to a priori enforce large heterogeneity in responses, I set σ_ξ^2 at 0.01. In the stacked vector $\widetilde{\beta}_s$ that was presented in Section 4.1, this information on σ_ξ^2 combined with no a priori expectation of correlation between the different effects results in a prior centering of Σ at $0.01 I_{(J+K)}$. Under an inverse Wishart prior for Σ, the minimum prior degrees of freedom that ensures a prior mean is $J + K + 3$, which I utilize here in order to remain diffuse about the degree of heterogeneity in the coefficients.

I also must specify prior information on the coefficients on the long-run deviation parameter λ_s. I maintain that it should be centered at -0.5, which implies that $w_s \lambda^*$ should be centered at -0.5. As with ξ^*, I achieve this by setting the coefficient on the constant to have a prior mean of -0.5, while the remaining coefficients are centered at 0. I also preserve the prior variance from the first section, so that under a normality assumption, $\lambda^* \sim \mathcal{N}\left(\underline{\lambda}^*, I_L\right)$, with $\underline{\lambda}^*$ as just described. Again, I do not want to introduce heterogeneity in the λ_s inadvertently through the prior specification, and so specify that $\sigma_\lambda^2 \sim \mathcal{IG}\left(\nu/2, 2/\nu s^2\right)$, with centrality parameter s^2 set at 0.01 and $\nu = 3$ to emphasize diffuseness.

An analogous prior distribution is specified for θ_s. Because I want $w_s \theta^*$ to be centered at 1 with small variance, I set the first element of $\underline{\theta}^*$ to be 1 and the remaining coefficients at 0, and then place a prior variance on each element of θ^* of 0.1/6, similar to the prior for the basic model. The prior information on σ_θ^2 is identical to that for σ_λ^2.

Last among the first-stage coefficients is the prior information on the long-run markup parameters $\widetilde{\alpha}_s$. I still expect an average markup of retail over spot price of about 80¢. As above, I set the first element of α^* to be 0.8 and the remaining parameters at 0, and each with a variance of 5. The spatial covariance matrix depends on the vector of parameters $\left(\tau_\alpha^2, \sigma_\alpha^2, \phi_a\right)$. These parameters must remain in the positive domain to ensure a stable and positive definite covariance matrix A. For this reason, I work with the reparameterized vector $\rho = \log\left(\tau_\alpha^2, \sigma_\alpha^2, \phi_a\right)$, about which I assume a normal prior centered at $\log(0.1, 0.1, 0.01)$ with prior covariance matrix I_3. With ϕ_α centered at 0.01, there is a strong a priori belief that the correlation in markups between stations depends inversely on the distance between them, although with the prior information on τ_α^2 and σ_α^2, the expected correlation is bounded above by 0.5.

Finally, with regard to the residual variance, I maintain the original specification from above. The only difference is that I expect the fit to improve and that the number of effective regressors has increased from $J+K+S+2$ to $S(J+K+3)$. I increase my prior expectation of R^2 from 0.25 to 0.75, in accordance with the expectation that adding each of the station effects will sharply improve fit relative to the basic ECM.

A.4 Posterior estimation and output of the hierarchical ECM

Chib and Carlin (1999) describe the basic algorithm for estimating a hierarchical model like the one employed here. I modify this algorithm to account for the nonlinear procedure that led to the Gibbs sampler in Section A.2. For all of the parameters except $\widetilde{\alpha}_s$ and ρ, because of the near-conjugacy of my prior this is a trivial extension to the basic Gibbs sampler. The only complication for $\widetilde{\alpha}_s$ is that unlike the other parameters in the model, the nonindependence between station effects implies that I must draw the entire vector of coefficients together. Letting

$$y_{st} = -\frac{\Delta p_{st} - x_{st}\beta_s - \lambda_s p_{s,t-1}}{\lambda_s} - \theta_s c_{t-1}$$

and

$$v_s = -u_s/\lambda_s,$$

then stacking up the T_s time observations for station s, we have

$$y_s = \iota_s \widetilde{\alpha}_s + v_s,$$

where ι_s denotes a T_s vector of ones. Further letting $\mathcal{I} = blockdiag\left(\{\iota_s\}\right),$ we obtain the full vector of stacked observations (first over time, then over station)

$$y = \mathcal{I}\widetilde{\alpha} + v.$$

Recall that $\widetilde{\alpha} \sim \mathcal{N}\left(W\alpha^*, A\right)$ and let $\mathcal{T} = diag\left(\left\{T_s\lambda_s^2\right\}\right)$ and $\overline{y} = \left(\overline{y}_1, ...\overline{y}_S\right)'$ equal the vector of sample averages over time of y_s for each station. Then following Lindley and Smith (1972), we have that $\widetilde{\alpha}|y, \sigma_u^2, \alpha^*, A \sim \mathcal{N}\left(Dd, D\right),$ where

$$D = \left(\mathcal{T}/\sigma_u^2 + A^{-1}\right)^{-1}$$

$$d = \mathcal{T}\overline{y}/\sigma_u^2 + A^{-1}W\alpha^*.$$

Given a draw for $\widetilde{\alpha}$, I obtain a draw for α^* in the usual way. But the Gibbs sampler cannot be used to obtain a draw for ρ, and so for this study I rely on the Metropolis Hastings algorithm with a multivariate t candidate density, where the posterior for ρ is proportional to

$$\phi\left(\rho; \underline{\rho}, \underline{V}_\rho\right)\phi\left(\widetilde{\alpha}; W\alpha^*, A\left(\rho\right)\right)$$

and the hyperparameters $\underline{\rho}$ and $\underline{V}_\rho$ are discussed in Section A.3 above.

Regarding final output of the results, I mentioned in the text that there are far too many parameters—403 coefficient parameters, the 10x10 matrix Σ, the two variance parameters associated with λ_s and θ_s, and the three covariance function parameters ρ in A—for me to present a summary of all results in this paper, although I can provide them to the interested reader upon request. A few parameters are worth discussing, however. In particular, I find a posterior mean for ϕ_α of 0.0039 with a posterior standard deviation of 0.003. Together with the other parameters in ρ, this implies that the correlation between stations that are 1 mile apart is 0.975, while the correlation between stations that are 5 miles apart is 0.961.

I find this to be surprising slow decay, and take it as evidence of the inappropriateness of an independence assumption in the markups.

Also, I continue to find, as with the basic ECM, that θ_s is estimated to be significantly greater than 1. In fact, the distribution over s of the posterior mean of each θ_s is centered at 1.36 with a standard deviation of 0.11, implying that for the average station, when costs rise by 10¢, its price rises by 13.6¢. However, there is considerable variation across stations: the station at the 25th percentile has an estimated θ_s of 1.27 while for the station at the 75th percentile this value is 1.43. One implication of this is that the markup, $p - c$, is not simply just the parameter $\widetilde{\alpha}$, but rather $\widetilde{\alpha} + (\theta - 1)\, c$. In Table A2, I summarize the predicted marginal effects of each of the station characteristics on the estimated markup of retail over wholesale spot prices when the spot price is $1.

Tables & Figures

Table 1

Descriptive Statistics (In Dollars)

	Average	St. Dev.	Min	Max	n
Retail Price	1.794	0.223	1.439	2.359	4466
Wholesale Price	0.944	0.160	0.719	1.503	46
Retail Price Change	0.006	0.054	-0.360	0.280	4109
Wholesale Price Change	0.002	0.066	-0.183	0.170	45
Pos. Retail Price Change	0.061	0.033	< 0.001	0.280	1448
Neg. Retail Price Change*	-0.026	0.035	-0.360	0	2661
Pos. Wholesale Price Change	0.042	0.038	0.001	0.170	27
Neg. Wholesale Price Change	-0.058	0.052	-0.183	-0.004	18

* Includes 1060 no-change observations.

Table 2

Descriptive Statistics of Station Characteristics

Variable	Mean	Std. Dev.	Min	Max
Arco	0.126	-	-	-
Chevron	0.227	-	-	-
Mobil	0.185	-	-	-
Shell	0.134	-	-	-
Texaco*	0.076	-	-	-
Unocal 76	0.160	-	-	-
Number Rivals < 0.1 Miles	0.437	0.630	0.000	2.0000
Number Rivals 0.1 to 0.5 Miles	1.092	1.390	0.000	7.0000
Number Rivals 0.5 to 1.0 Miles	1.345	1.429	0.000	7.0000
Distance to Nearest Rival	0.413	0.530	0.029	2.6155
Independent Owned	0.286	-	-	-
Jobber Owned	0.084	-	-	-
Major owned: Lessee-Dealer	0.403	-	-	-
Major Owned: Salary Operated	0.227	-	-	-
Distance from Freeway	2.121	2.436	0.016	8.910
Pumps per Acre	24.147	10.605	8.534	61.952
Lot Size	0.467	0.181	0.129	0.918
Carwash	0.294	-	-	-
Service Station	0.252	-	-	-
Convenience Store	0.504	-	-	-
Island Kiosk	0.034	-	-	-
Visible Competitor Prices	0.412	-	-	-
Visible Freeway Sign	0.076	-	-	-
No. Driveways	2.311	0.828	1.000	4.000
Shopping Center	0.479	-	-	-
Traffic Light	0.101	-	-	-
Divided Primary Arterial	0.496	-	-	-
Nearby Residential	0.496	-	-	-
Nearby Commercial	0.412	-	-	-
Population Size	1.503	0.285	0.936	2.086
Population Density	6.527	1.877	0.337	11.180
Housing Density	2.738	1.031	0.184	4.768
Percent Commuting < 5 Miles	2.801	3.251	0.812	23.867
Median Household Income	73.668	14.578	39.430	108.470
Median Rent	1.127	0.143	0.885	1.455

*From November 2002 to January 2003, all of the Texaco Stations were switched to Shell.

Table 3

Mean Predictive Effects of Covariates on Rate of Return to Long-Run Equilibrium

Variable	Posterior Distribution Mean	Std. Dev.	Pr(> 0)
Constant	-0.176	0.134	0.088
Arco	0.063	0.039	0.945
Chevron	-0.021	0.037	0.281
Mobil	0.031	0.040	0.795
Shell	-0.023	0.041	0.284
Texaco	-0.065	0.049	0.093
Unocal 76	-0.002	0.039	0.486
Number Rivals < 0.1 Miles	-0.005	0.018	0.374
Number Rivals 0.1 to 0.5 Miles	0.015	0.006	0.995
Number Rivals .5 to 1.0 Miles	0.005	0.007	0.755
Distance to Nearest Rival	0.010	0.019	0.687
Independent Owned	0.005	0.037	0.539
Major owned: Lessee-Dealer	-0.009	0.042	0.393
Major Owned: Salary Operated	-0.008	0.039	0.400
Distance from Freeway	-0.001	0.004	0.464
Pumps per Acre	0.000	0.001	0.553
Lot Size	-0.001	0.067	0.495
Carwash	0.044	0.023	0.965
Service Station	0.016	0.020	0.793
Convenience Store	0.010	0.017	0.712
Island Kiosk	-0.065	0.053	0.110
Visible Competitor Prices	0.031	0.027	0.871
Visible Freeway Sign	0.017	0.034	0.691
No. Driveways	-0.007	0.013	0.271
Shopping Center	-0.001	0.020	0.508
Traffic Light	-0.055	0.035	0.049
Divided Primary Arterial	-0.003	0.018	0.428
Nearby Residential	-0.003	0.020	0.459
Nearby Commercial	-0.022	0.018	0.116
Population Size	-0.018	0.035	0.313
Population Density	-0.009	0.006	0.082
Percent Commuting < 5 Miles	-0.008	0.004	0.014
Median Household Income	0.000	0.001	0.579

Table A1

Posterior Distribution of Regression Parameters

Dependent Variable: $\Delta Retail_t$ Covariate:	Posterior Mean	St. Dev.	$\Pr(\cdot > 0 \mid Data)$
$\Delta Wholesale_t^-$ (β_0^-)	0.198	0.021	> 0.999
$\Delta Wholesale_t^+$ (β_0^+)	0.080	0.036	0.988
$\Delta Wholesale_{t-1}^-$ (β_1^-)	0.005	0.023	0.590
$\Delta Wholesale_{t-1}^+$ (β_1^+)	-0.165	0.038	< 0.001
$\Delta Wholesale_{t-2}^-$ (β_2^-)	-0.245	0.020	< 0.001
$\Delta Wholesale_{t-2}^+$ (β_2^+)	0.392	0.036	> 0.999
$\Delta Retail_{t-1}^-$ (γ_1^-)	-0.136	0.022	< 0.001
$\Delta Retail_{t-1}^+$ (γ_1^+)	0.285	0.036	> 0.999
$\Delta Retail_{t-2}^-$ (γ_2^-)	-0.218	0.023	< 0.001
$\Delta Retail_{t-2}^+$ (γ_2^+)	0.267	0.035	> 0.999
Long Run Cost Response (θ)	1.376	0.023	> 0.999
Long Run Deviation (λ)	-0.265	0.007	< 0.001
Regression Std.Dev.	0.00109	0.00002	1

Table A2

Mean Predictive Effects of Covariates on Markups (Retail - Spot Price)

	Posterior Distribution		
Variable	Mean	Std. Dev.	Pr(> 0)
Constant	0.756	0.2301	> 0.999
Arco	0.011	0.034	0.631
Chevron	0.061	0.028	0.982
Mobil	0.053	0.033	0.937
Shell	0.031	0.034	0.816
Texaco	0.033	0.034	0.832
Unocal 76	0.065	0.032	0.976
Number Rivals < 0.1 Miles	0.026	0.020	0.903
Number Rivals 0.1 to 0.5 Miles	0.002	0.007	0.626
Number Rivals .5 to 1.0 Miles	0.005	0.007	0.760
Distance to Nearest Rival	0.001	0.019	0.531
Independent Owned	0.004	0.031	0.553
Major owned: Lessee-Dealer	0.003	0.036	0.531
Major Owned: Salary Operated	-0.010	0.033	0.371
Distance from Freeway	0.003	0.005	0.691
Pumps per Acre	-0.002	0.001	0.009
Lot Size	-0.086	0.059	0.072
Carwash	0.019	0.018	0.858
Service Station	-0.006	0.016	0.356
Convenience Store	0.007	0.014	0.693
Island Kiosk	-0.050	0.042	0.120
Visible Competitor Prices	-0.016	0.025	0.263
Visible Freeway Sign	0.015	0.034	0.676
No. Driveways	0.002	0.010	0.576
Shopping Center	0.006	0.015	0.658
Traffic Light	0.012	0.029	0.668
Divided Primary Arterial	-0.002	0.015	0.447
Nearby Residential	-0.004	0.016	0.396
Nearby Commercial	-0.001	0.015	0.480
Population Size	0.001	0.033	0.513
Population Density	-0.003	0.006	0.332
Percent Commuting < 5 Miles	-0.003	0.003	0.193
Median Household Income	0.001	0.001	0.930

Table A3

Posterior Distribution of Regression Parameters

Variable	Δc_t			$\Delta c_t > 0$		
	Mean	Std. Dev	Pr(> 0)	Mean	Std. Dev	Pr(> 0)
Constant	0.5891	0.3119	0.9678	-1.0901	0.5631	0.0458
Arco	-0.1386	0.0912	0.0936	0.1797	0.1261	0.9116
Chevron	-0.1063	0.0696	0.0826	0.1475	0.1066	0.9188
Mobil	-0.1479	0.0988	0.0802	0.4006	0.1560	0.9788
Shell	-0.0802	0.1141	0.2136	0.1035	0.1687	0.7426
Texaco	0.0432	0.0993	0.6376	-0.1125	0.1488	0.2656
Unocal	-0.0820	0.1035	0.2158	0.2274	0.1568	0.9016
Number Rivals < 0.1 Miles	0.0527	0.0449	0.8738	-0.0936	0.0845	0.1536
Number Rivals0.1 to 0.5 Miles	0.0217	0.0181	0.8492	-0.0421	0.0327	0.1422
Number Rivals 0.5 to 1.0 Miles	0.0159	0.0201	0.7332	-0.0359	0.0323	0.1778
Distance to Nearest Rival	0.0076	0.0452	0.5590	0.0232	0.0836	0.6628
Independent Owned	-0.0257	0.0884	0.3388	0.0195	0.1466	0.6018
Major owned: Lessee-Dealer	-0.0015	0.1088	0.4586	0.0490	0.1874	0.5598
Major owned: Salary Operated	-0.0097	0.1038	0.4670	0.0933	0.1914	0.6708
Distance from Freeway	0.0117	0.0084	0.9038	-0.0199	0.0140	0.0784
Pumps Per Acre	-0.0072	0.0023	0.0052	0.0140	0.0038	1.0000
Lot Size	-0.0102	0.1631	0.5402	0.1991	0.2696	0.7488
Carwash	0.0409	0.0419	0.8244	-0.0899	0.0756	0.1044
Service Station	0.0188	0.0384	0.7018	-0.0353	0.0705	0.3248
Convenience Store	0.0111	0.0388	0.5984	0.0031	0.0714	0.5682
Island Kiosk	0.4730	0.0861	1.0000	-0.8103	0.1286	0.0000
Visible Competitor Prices	-0.1422	0.0720	0.0028	0.2623	0.1152	0.9956
Visible Freeway Sign	-0.0532	0.1033	0.3190	0.1116	0.1593	0.7758
No. Driveways	0.0251	0.0293	0.8282	-0.0207	0.0547	0.3946
Shopping Center	0.0045	0.0343	0.5288	0.0890	0.0606	0.9210
Traffic Light	0.0878	0.0673	0.9046	-0.0282	0.1225	0.4244
Divided Primary Arterial	0.0326	0.0471	0.7352	-0.1326	0.0771	0.0290
Nearby Residential	0.0773	0.0588	0.8934	-0.1092	0.0996	0.1220
Nearby Commercial	-0.0006	0.0363	0.5126	0.0533	0.0686	0.7530
Population Size	-0.0620	0.0650	0.1848	0.1185	0.1218	0.8040
Population Density	-0.0175	0.0169	0.1528	0.0411	0.0257	0.9396
Percent Commuting < 5 Miles	-0.0128	0.0059	0.0082	0.0277	0.0093	1.0000
Median Household Income	-0.0007	0.0022	0.3486	0.0024	0.0038	0.7474
Standard Deviation	0.0164	0.0018	1.0000	0.0168	0.0019	1.0000

Table A3

Posterior Distribution of Regression Parameters

Variable	Δc_{t-1}			$\Delta c_{t-1} > 0$		
	Mean	Std. Dev	Pr(> 0)	Mean	Std. Dev	Pr(> 0)
Constant	-0.4356	0.3775	0.1026	1.1344	0.4799	1.0000
Arco	0.1590	0.0971	0.9580	-0.0557	0.1501	0.3756
Chevron	0.0605	0.0851	0.7370	-0.2016	0.1345	0.0660
Mobil	0.2851	0.0729	1.0000	-0.6174	0.1108	0.0000
Shell	0.1232	0.1083	0.8674	-0.1475	0.1635	0.1846
Texaco	0.0247	0.1758	0.6374	0.1058	0.2380	0.5702
Unocal	0.1206	0.0879	0.9176	-0.2481	0.1537	0.0424
Number Rivals < 0.1 Miles	-0.0076	0.0461	0.4600	0.0210	0.0699	0.6032
Number Rivals0.1 to 0.5 Miles	0.0194	0.0146	0.9034	0.0188	0.0243	0.7910
Number Rivals 0.5 to 1.0 Miles	0.0094	0.0228	0.5966	0.0159	0.0364	0.6580
Distance to Nearest Rival	0.0029	0.0462	0.5444	0.0368	0.0667	0.7138
Independent Owned	-0.0183	0.1020	0.4246	0.0617	0.1227	0.6756
Major owned: Lessee-Dealer	-0.0446	0.1309	0.3812	0.0760	0.1663	0.6428
Major owned: Salary Operated	-0.0208	0.1319	0.4214	-0.0457	0.1656	0.4328
Distance from Freeway	-0.0238	0.0100	0.0136	0.0366	0.0190	0.9880
Pumps Per Acre	0.0024	0.0027	0.8216	-0.0055	0.0039	0.0672
Lot Size	-0.1535	0.1831	0.1880	0.1813	0.2605	0.7830
Carwash	0.1097	0.0622	0.9496	-0.0359	0.0887	0.4080
Service Station	0.0265	0.0547	0.6522	0.0508	0.0728	0.7124
Convenience Store	0.0579	0.0514	0.8524	-0.0603	0.0702	0.2222
Island Kiosk	-0.2411	0.1427	0.0520	0.5598	0.2074	1.0000
Visible Competitor Prices	0.1003	0.0987	0.7984	-0.1139	0.1263	0.2066
Visible Freeway Sign	-0.0359	0.0953	0.3410	0.0029	0.1696	0.5084
No. Driveways	0.0135	0.0292	0.7024	-0.0297	0.0416	0.1824
Shopping Center	-0.0687	0.0462	0.0948	0.0579	0.0897	0.7420
Traffic Light	-0.0745	0.0901	0.2096	-0.0201	0.1005	0.4180
Divided Primary Arterial	0.0253	0.0546	0.6610	0.0132	0.0831	0.5682
Nearby Residential	-0.0501	0.0507	0.1606	0.0941	0.0785	0.8814
Nearby Commercial	0.0012	0.0444	0.5172	-0.0997	0.0729	0.0532
Population Size	0.0586	0.0846	0.7586	-0.1773	0.1223	0.0822
Population Density	0.0001	0.0149	0.4840	-0.0398	0.0261	0.0972
Percent Commuting < 5 Miles	-0.0061	0.0074	0.2064	-0.0171	0.0113	0.0552
Median Household Income	0.0035	0.0031	0.8764	-0.0080	0.0041	0.0512
Standard Deviation	0.0171	0.0021	1.0000	0.0175	0.0022	1.0000

Table A3

Posterior Distribution of Regression Parameters

Variable	Δc_{t-2}			$\Delta c_{t-2} > 0$		
	Mean	Std. Dev	Pr(> 0)	Mean	Std. Dev	Pr(> 0)
Constant	-0.1340	0.3319	0.3740	0.3781	0.6236	0.6846
Arco	-0.0001	0.1098	0.4624	0.0441	0.2039	0.5914
Chevron	-0.0404	0.0829	0.3204	0.1409	0.1376	0.8376
Mobil	-0.0788	0.1131	0.2544	0.3131	0.2083	0.9580
Shell	-0.0970	0.1106	0.2380	0.2577	0.1938	0.9230
Texaco	-0.2177	0.1255	0.0422	0.2745	0.2349	0.8796
Unocal	-0.0588	0.1109	0.2898	0.1649	0.2034	0.7634
Number Rivals < 0.1 Miles	0.0111	0.0526	0.5376	0.0615	0.1011	0.7518
Number Rivals0.1 to 0.5 Miles	0.0114	0.0128	0.8050	0.0127	0.0272	0.7414
Number Rivals 0.5 to 1.0 Miles	-0.0108	0.0131	0.2130	0.0359	0.0244	0.9580
Distance to Nearest Rival	0.0283	0.0426	0.7812	-0.0434	0.0688	0.2514
Independent Owned	0.0687	0.1064	0.7180	-0.1705	0.1999	0.1920
Major owned: Lessee-Dealer	0.0341	0.1077	0.6154	-0.0917	0.1997	0.3866
Major owned: Salary Operated	0.0628	0.1165	0.6720	-0.0736	0.2432	0.4658
Distance from Freeway	0.0141	0.0081	0.9504	-0.0282	0.0153	0.0372
Pumps Per Acre	0.0027	0.0027	0.8744	-0.0052	0.0052	0.1520
Lot Size	0.2269	0.1793	0.8806	-0.3461	0.3728	0.2412
Carwash	-0.0438	0.0457	0.1766	0.1640	0.0868	0.9876
Service Station	-0.0380	0.0540	0.2536	0.0944	0.1125	0.7958
Convenience Store	-0.0157	0.0457	0.3596	0.0883	0.0967	0.7924
Island Kiosk	-0.2836	0.1263	0.0288	0.4686	0.1946	0.9858
Visible Competitor Prices	0.0382	0.0778	0.6634	-0.0878	0.1519	0.3102
Visible Freeway Sign	0.0494	0.0773	0.7196	-0.1044	0.1264	0.2024
No. Driveways	-0.0424	0.0275	0.0712	0.0771	0.0501	0.9442
Shopping Center	0.0795	0.0383	0.9860	-0.1407	0.0727	0.0338
Traffic Light	-0.0438	0.0721	0.2622	0.0570	0.1208	0.6794
Divided Primary Arterial	-0.0184	0.0368	0.3310	0.0657	0.0619	0.8336
Nearby Residential	-0.0037	0.0597	0.4372	0.0321	0.0963	0.6526
Nearby Commercial	-0.0083	0.0330	0.4004	-0.0100	0.0581	0.3912
Population Size	-0.0332	0.0778	0.2972	0.0135	0.1538	0.5636
Population Density	-0.0096	0.0183	0.2772	-0.0091	0.0293	0.4470
Percent Commuting < 5 Miles	-0.0048	0.0056	0.1976	-0.0091	0.0089	0.1616
Median Household Income	-0.0011	0.0031	0.3812	0.0007	0.0060	0.5488
Standard Deviation	0.0178	0.0023	1.0000	0.0182	0.0025	1.0000

Table A3

Posterior Distribution of Regression Parameters

Variable	Δp_{t-1}			$\Delta p_{t-1} > 0$		
	Mean	Std. Dev	Pr(> 0)	Mean	Std. Dev	Pr(> 0)
Constant	-0.4486	0.5450	0.2618	0.7746	0.8294	0.7674
Arco	0.1657	0.1242	0.8910	-0.1372	0.1517	0.1610
Chevron	-0.3391	0.1214	0.0000	0.5203	0.1494	1.0000
Mobil	-0.0933	0.1261	0.2148	0.2119	0.1467	0.8876
Shell	-0.1641	0.1131	0.0824	0.1187	0.1173	0.8458
Texaco	-0.1654	0.1630	0.2008	0.0052	0.2584	0.5370
Unocal	-0.1256	0.1245	0.1978	-0.0055	0.1663	0.5526
Number Rivals < 0.1 Miles	-0.0034	0.0868	0.4422	-0.0367	0.1247	0.4288
Number Rivals0.1 to 0.5 Miles	0.0620	0.0231	1.0000	-0.0932	0.0369	0.0000
Number Rivals 0.5 to 1.0 Miles	-0.0113	0.0175	0.2962	-0.0264	0.0305	0.1902
Distance to Nearest Rival	0.1087	0.0445	1.0000	-0.1256	0.0681	0.0206
Independent Owned	0.0188	0.1428	0.5424	-0.0595	0.2228	0.3978
Major owned: Lessee-Dealer	-0.0153	0.1523	0.4234	-0.0484	0.2394	0.4608
Major owned: Salary Operated	0.0031	0.1525	0.4764	0.0286	0.2396	0.4852
Distance from Freeway	-0.0145	0.0131	0.1248	-0.0030	0.0202	0.3772
Pumps Per Acre	0.0019	0.0034	0.6998	0.0041	0.0052	0.8140
Lot Size	0.0064	0.2351	0.5548	0.1923	0.3274	0.7136
Carwash	0.1783	0.0713	1.0000	-0.3603	0.1150	0.0000
Service Station	0.0745	0.0527	0.9204	-0.1136	0.0829	0.1166
Convenience Store	0.0630	0.0425	0.9244	-0.1252	0.0586	0.0002
Island Kiosk	0.0698	0.1475	0.6612	-0.3662	0.1851	0.0104
Visible Competitor Prices	0.0505	0.0912	0.6724	-0.0593	0.1219	0.3736
Visible Freeway Sign	0.0488	0.1090	0.6240	0.0403	0.1845	0.5494
No. Driveways	0.0411	0.0312	0.9422	-0.1004	0.0535	0.0014
Shopping Center	-0.0394	0.0501	0.2140	0.0512	0.0720	0.7366
Traffic Light	0.0199	0.1146	0.6032	-0.1116	0.2033	0.3226
Divided Primary Arterial	-0.0701	0.0647	0.1510	0.0201	0.0768	0.5444
Nearby Residential	0.0301	0.0606	0.6886	-0.0708	0.0879	0.2264
Nearby Commercial	0.0706	0.0500	0.9358	-0.0861	0.0813	0.1554
Population Size	0.1476	0.1191	0.9096	-0.2792	0.1597	0.0256
Population Density	-0.0163	0.0296	0.3390	0.0452	0.0528	0.7486
Percent Commuting < 5 Miles	-0.0203	0.0089	0.0190	0.0409	0.0119	1.0000
Median Household Income	0.0015	0.0058	0.5316	-0.0009	0.0092	0.4850
Standard Deviation	0.0185	0.0026	1.0000	0.0189	0.0027	1.0000

Table A3

Posterior Distribution of Regression Parameters

Variable	Δp_{t-2}			$\Delta p_{t-2} > 0$		
	Mean	Std. Dev	Pr(> 0)	Mean	Std. Dev	Pr(> 0)
Constant	-0.5755	0.4352	0.0978	0.9850	0.5591	0.9480
Arco	0.1407	0.1168	0.9070	-0.3163	0.1647	0.0124
Chevron	0.1170	0.0956	0.8868	-0.3182	0.1189	0.0018
Mobil	0.1122	0.1260	0.7992	-0.2487	0.1387	0.0708
Shell	0.1523	0.1318	0.8574	-0.2343	0.1991	0.1366
Texaco	0.0173	0.1347	0.5156	0.0752	0.1536	0.7662
Unocal	0.0977	0.0935	0.8516	-0.1019	0.1099	0.2162
Number Rivals < 0.1 Miles	0.1247	0.0781	0.9448	-0.2060	0.1174	0.0474
Number Rivals0.1 to 0.5 Miles	0.0333	0.0148	0.9904	-0.0195	0.0208	0.1880
Number Rivals 0.5 to 1.0 Miles	0.0395	0.0222	0.9518	-0.0320	0.0295	0.1312
Distance to Nearest Rival	0.0652	0.0594	0.8642	-0.0513	0.0937	0.3130
Independent Owned	0.0597	0.1290	0.6834	0.0528	0.1950	0.5496
Major owned: Lessee-Dealer	0.0214	0.1507	0.5248	0.0634	0.2307	0.6156
Major owned: Salary Operated	0.0807	0.1182	0.7372	-0.0707	0.1968	0.4244
Distance from Freeway	0.0094	0.0124	0.7872	0.0022	0.0139	0.5816
Pumps Per Acre	0.0008	0.0024	0.6410	-0.0033	0.0044	0.2344
Lot Size	-0.1365	0.1809	0.2492	-0.0567	0.2389	0.4124
Carwash	0.0511	0.0769	0.7452	0.0367	0.1325	0.5412
Service Station	-0.0556	0.0743	0.2470	0.0587	0.1310	0.6228
Convenience Store	-0.0084	0.0475	0.4308	-0.0266	0.0488	0.2956
Island Kiosk	-0.2670	0.1415	0.0302	0.5074	0.2183	0.9934
Visible Competitor Prices	0.0288	0.0914	0.6184	0.0335	0.1279	0.6798
Visible Freeway Sign	0.0648	0.0906	0.7718	-0.0948	0.1166	0.2030
No. Driveways	0.0276	0.0363	0.7308	0.0218	0.0458	0.6606
Shopping Center	0.0865	0.0656	0.8862	-0.1579	0.0879	0.0736
Traffic Light	-0.1461	0.1014	0.0602	0.1577	0.1203	0.8990
Divided Primary Arterial	0.0522	0.0517	0.8214	0.0581	0.0596	0.8450
Nearby Residential	-0.0944	0.0603	0.0556	0.1274	0.1092	0.8594
Nearby Commercial	-0.0433	0.0616	0.2264	0.0663	0.0733	0.8096
Population Size	0.0517	0.1082	0.6416	-0.0114	0.1855	0.4992
Population Density	-0.0109	0.0241	0.3140	-0.0119	0.0297	0.4184
Percent Commuting < 5 Miles	-0.0109	0.0137	0.2050	0.0066	0.0179	0.7138
Median Household Income	0.0006	0.0024	0.6204	-0.0050	0.0035	0.0732
Standard Deviation	0.0192	0.0029	1.0000	0.0196	0.0030	1.0000

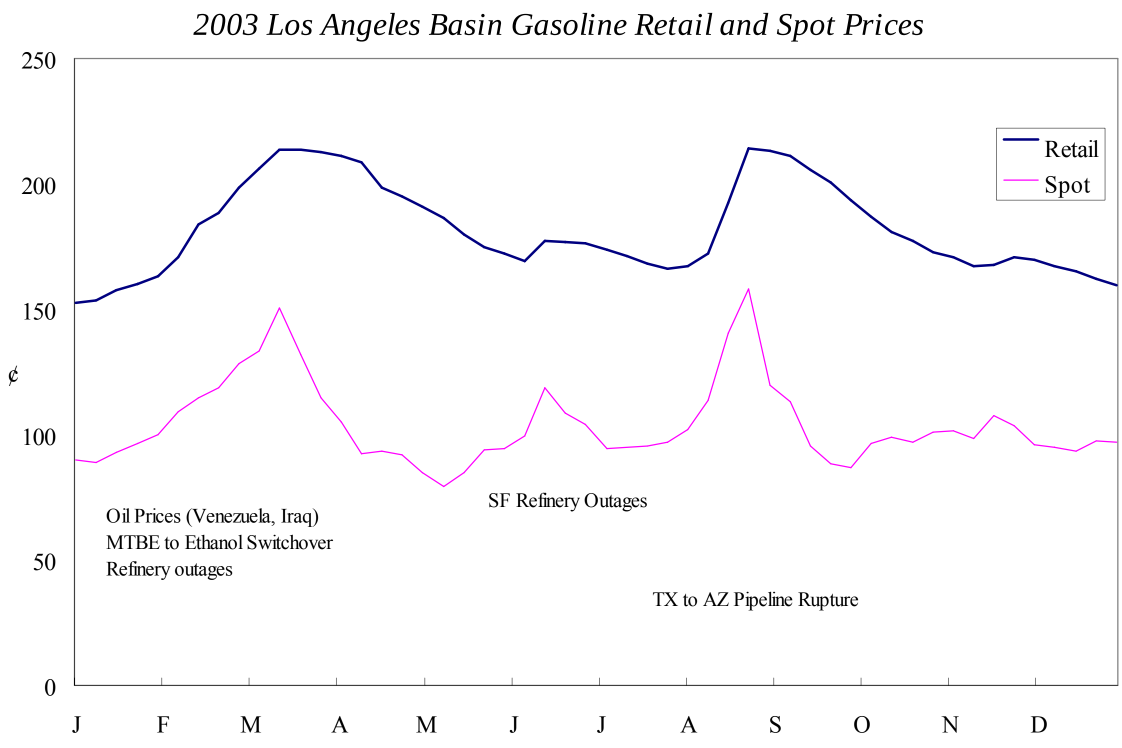

Figure 1: Plot of retail and wholesale spot prices for gasoline in Los Angeles.

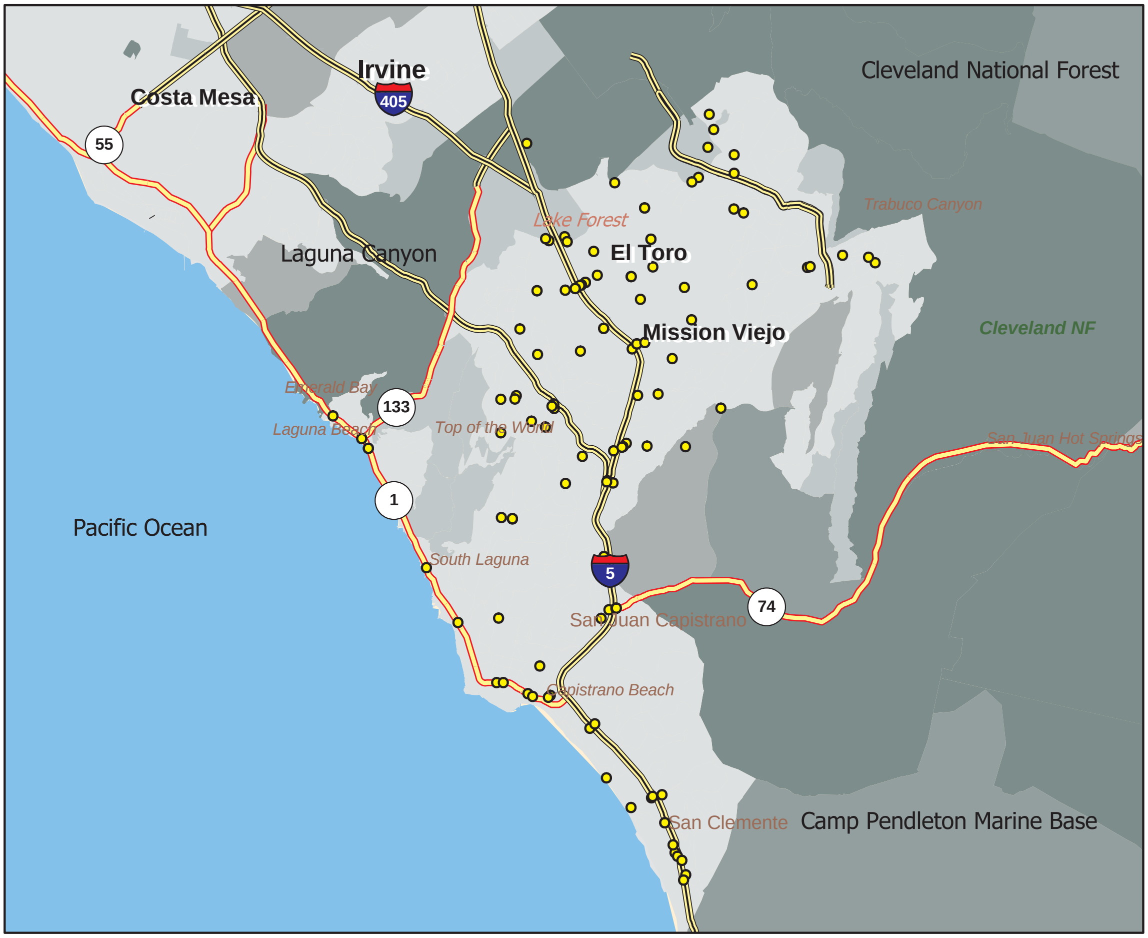

Figure 2: Map of South Orange County and Gasoline Stations.

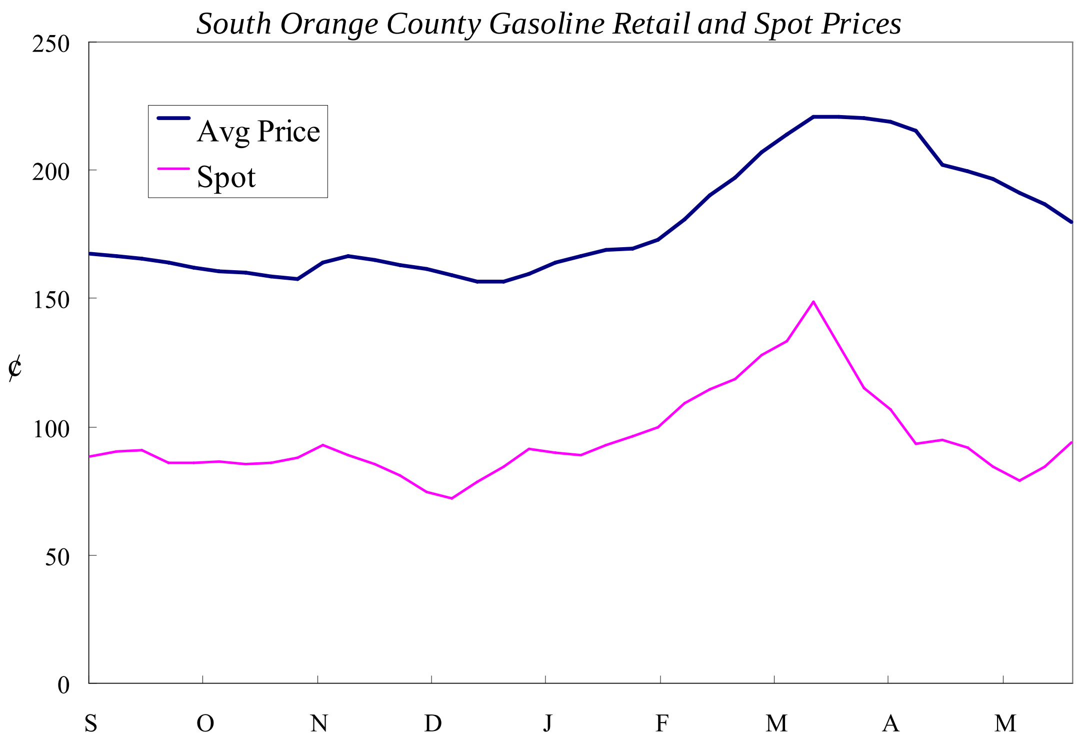

Figure 3: Plot of retail and wholesale spot prices for gasoline in South Orange County.

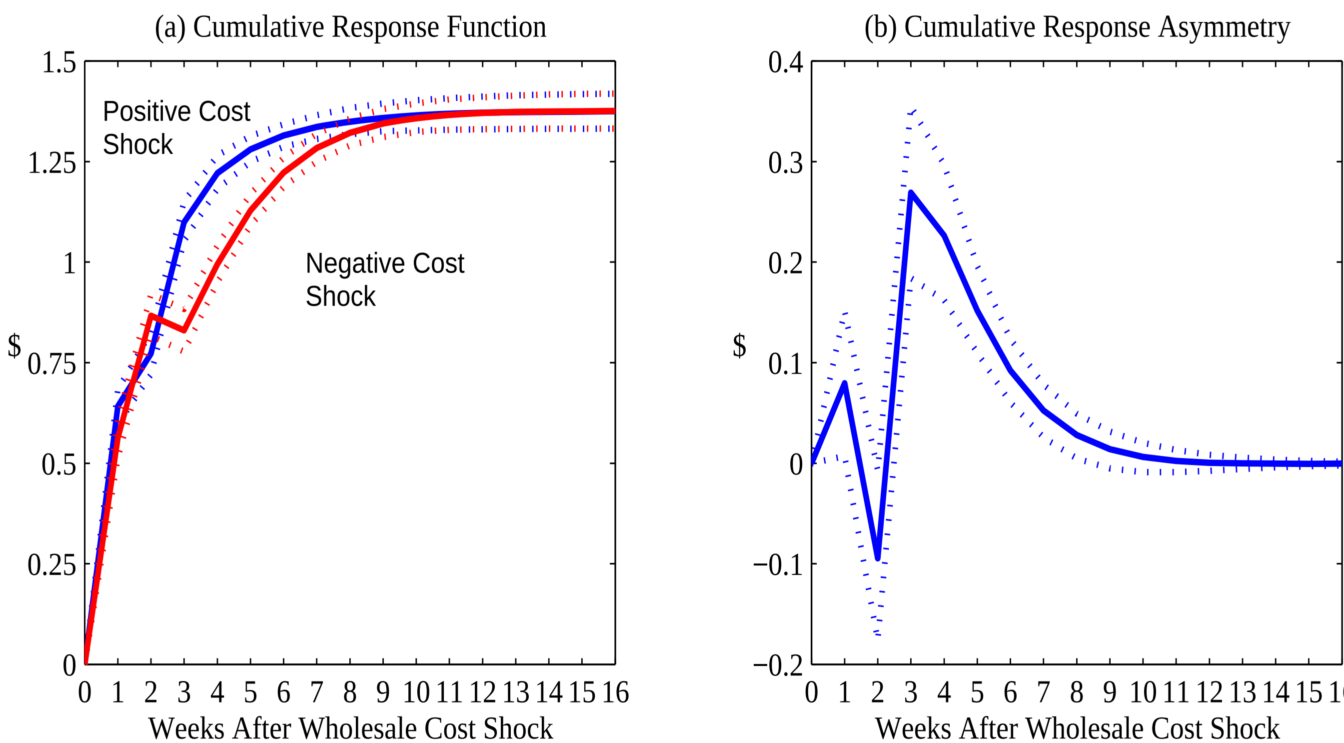

Figure 4: Plotting the effect of a $1 cost shock. Panel a describes the path of cumulative price changes for both a postive and a negative cost shock (modified to the positive domain). Panel b describes the difference in these two effects.

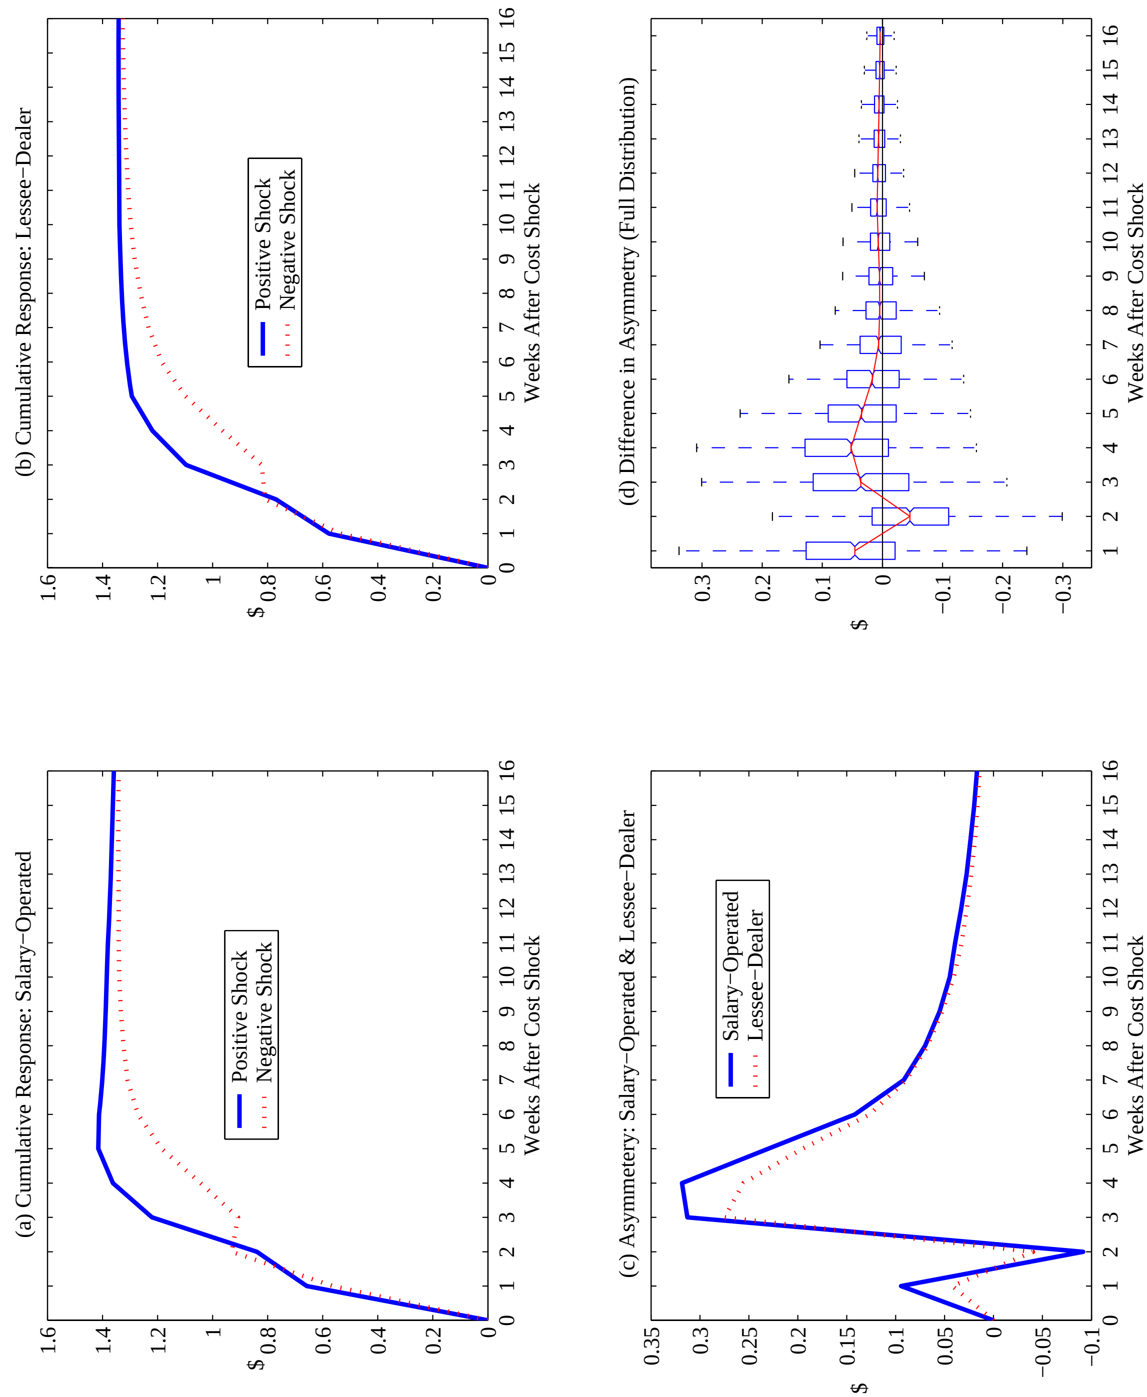

Figure 5: Charting the effect of a $1 cost shock for different organization types: Salary-operated versus lessee-dealer. Panel a describes the path of cumulative price changes for a salary-operated station. Panel b describes the path if these stations were switched to lessee-dealer. Panel c compares the difference in positive versus negative cost shocks for both types, and panel d describes the difference between the two effects from panel c.

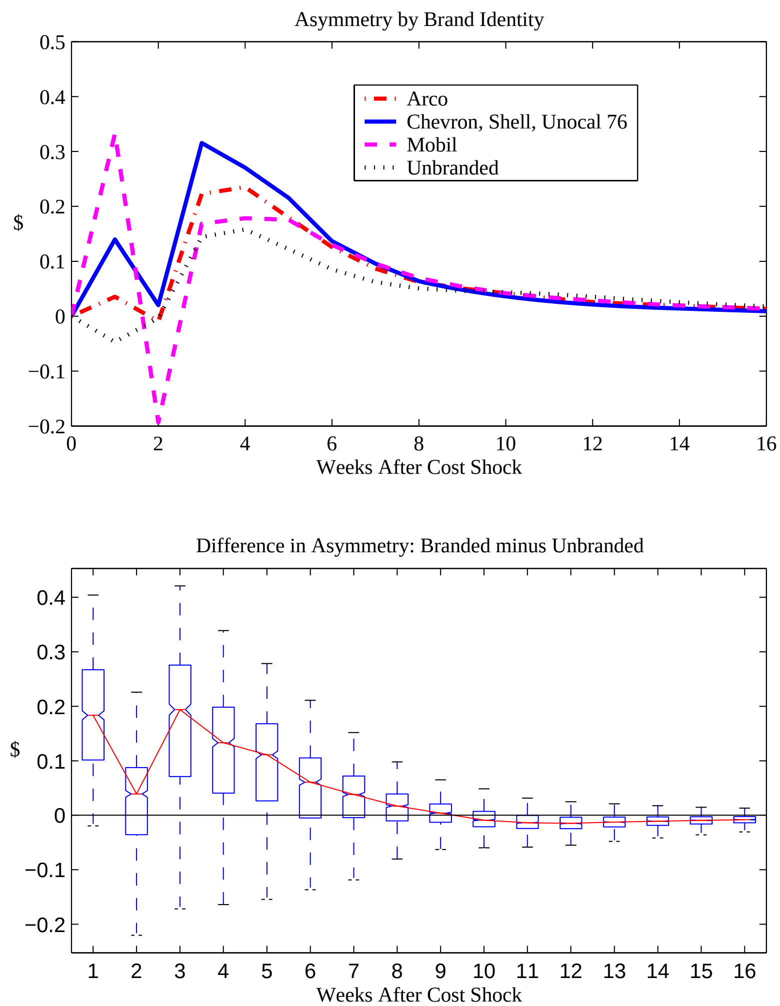

Figure 6: Charting the effect of brand identity on price response asymmetry. Panel a describes the difference in price responses for positive versus negative cost shocks across different brand types. Panel b plots the difference in asymmetries for branded versus unbranded stations.

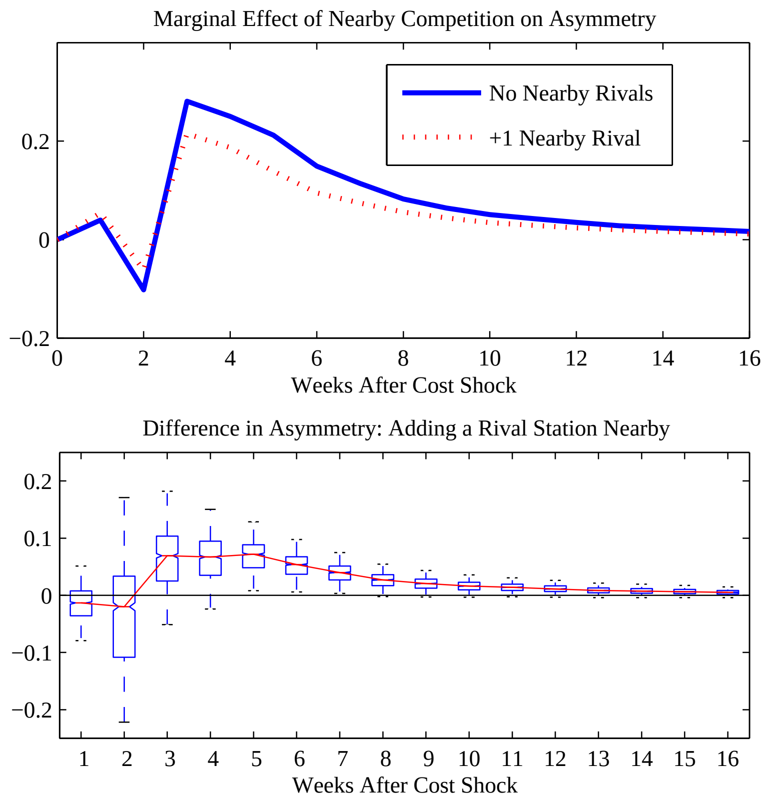

Figure 7: Charting the effect of rival proximity on price response asymmetry. Panel a describes the difference in price responses for positive versus negative cost shocks for stations with no immediate rivals and for stations with 1 rival added less than 0.1 miles away. Panel b plots the difference in these predictive asymmetries.

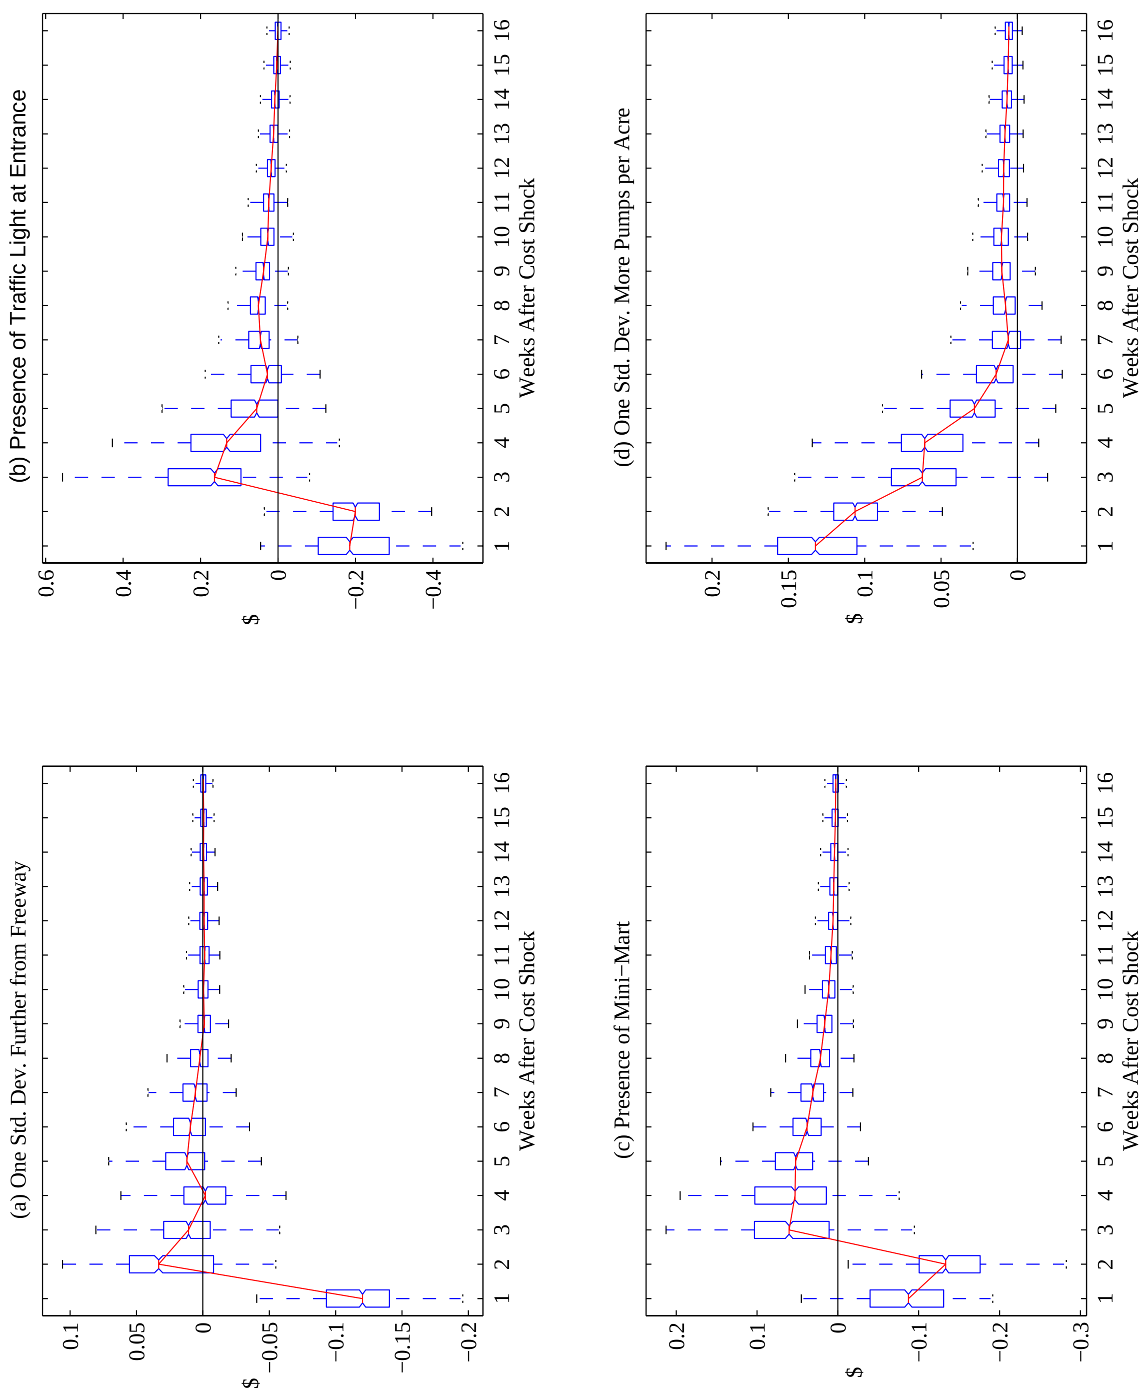

Figure 8: Charting the effects of site-level characteristics on price response asymmetry. Panels a and d describe the predictive change in price response asymmetry from increasing the observed characteristic for each station by one standard deviation. Panels b and c describe the predictive change for stations in having the observed characteristic versus removing it.

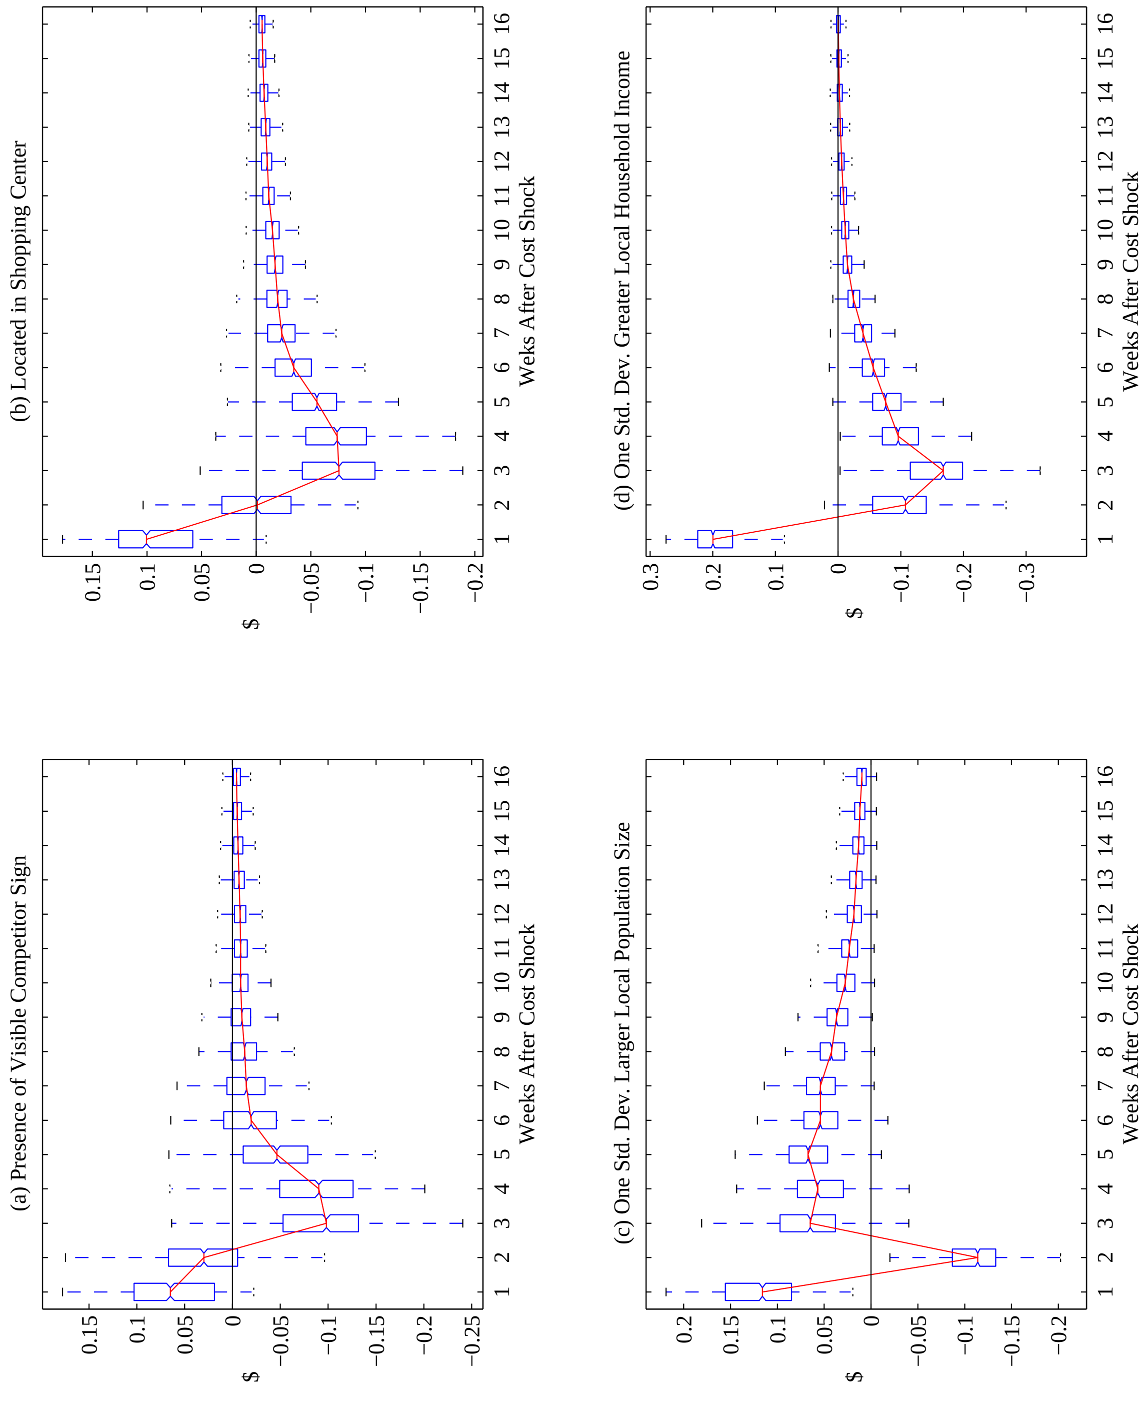

Figure 9: Charting the effects of local market-level characteristics on price response asymmetry. Panels a and b describe the predictive change in price response asymmetry for stations in having the observed characteristic versus removing it. Panels c and d describe the predictive change from increasing the observed characteristic for each station by one standard deviation.

References

Balke, Nathan S., Stephen P.A. Brown, and Mine K. Yücel, "Crude Oil and Gasoline Prices: An Asymmetric Relationship?" Federal Reserve Bank of Dallas Economic Review (First Quarter 1998).

Banerjee, Sudipto, Bradley P. Carlin, and Alan E. Gelfand, "Hierarchical Modeling and Analysis for Spatial Data" Chapman & Hall, Boca Raton (2004).

Barron, John M., John R. Umbeck, and Glen R. Waddell, "Consumer and Competitor Reactions: Evidence From a Retail-Gasoline Field Experiment" Working Paper, Department of Economics, Purdue University and Department of Economics, University of Oregon (September 2003).

Bauwens, Luc, Michel Lubrano, and Jean-François Richard, "Bayesian Inference in Dynamic Econometric Models" Oxford University Press, Oxford (1999).

Borenstein, Severin, "Selling Costs and Switching Costs: Explaining Retail Gasoline Margins," RAND Journal of Economics 22 (Autumn 1991), 354-369.

Borenstein, Severin and Andrea Shepard, "Dynamic Pricing in Retail Gasoline Markets" RAND Journal of Economics 27 (1996), 429-451.

Borenstein, Severin and Andrea Shepard, "Sticky Prices, Inventories, and Market Power in Wholesale Gasoline Markets" RAND Journal of Economics 33 (2002), 116-139.

Borenstein, Severin, A. Colin Cameron, and Richard Gilbert, "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?" Quarterly Journal of Economics 112 (February 1997), 305-339.

Bulow, Jeremy I. and Paul Pfleiderer, "A Note on the Effect of Cost Changes on Prices," Journal of Political Economy 91 (February 1983), 182-185.

Carlton, Dennis W. and Jeffrey M. Perloff, "Modern Industrial Organization" 4th Ed., Addison Wesley, Boston (2005).

Chang, Myong-Hun, "The Effects of Product Differentiation on Collusive Pricing" International Journal of Industrial Organization 9 (1991), 453-469.

Chang, Myong-Hun, "Intertemporal Product Choice and its Effects on Collusive Firm Behavior" International Economic Review 4 (1992), 773-793.

Chib, Siddhartha and Bradley P. Carlin, "On MCMC Sampling in Hierarchical Longitudinal Models" Statistics and Computing 9 (1999), 17-26.

Dahl, Carol and Thomas Sterner, "Analysing Gasoline Demand Elasticities: A Survey" Energy Economics 13 (1991), 203-210.

Deltas, George, "Retail Gasoline Price Response Stickiness and Asymmetry to Wholesale Price Changes: The Effect of Local Market Power" Working Paper, Department of Economics, University of Illinois-Urbana Champaign (August 2004).

Duffy-Deno, Kevin T., "Retail Price Asymmetries in Local Gasoline Markets" Energy Economics 18 (1996), 81-92.

(EIA) Energy Information Administration "2003 Gasoline Price Study: Final Report" SR/O&G/2003-03 (November 2003).

Engle, Robert F. and C. W.J. Granger, "Co-Integration and Error Correction: Representation, Estimation, and Testing" Econometrica 55 (March 1987), 251-276.

Gupta, Barnali and Guhan Venkatu, "Tacit Collusion in a Spatial Model with Delivered Pricing" Journal of Economics 76 (2002), 49-64.

Hackner, Jonas, "Endogenous product design in an infinitely repeated game" International Journal of Industrial Organization 13 (1995), 277-299.

Hackner, Jonas, "Optimal Symmetric Punishments in a Bertrand Differentiated Products Duopoly" International Journal of Industrial Organization 14 (1996), 611-630.

Haltiwanger, John. and Joseph E. Harrington Jr., "The Impact of Cyclical Demand Movements in Collusive Behavior" RAND Journal of Economics 22 (1991), 89-106.

Hastings, Justine, "Vertical Relationships and Competition in Retail Gasoline Markets: An Empirical Evidence from Contract Changes in Southern California" Competition Policy Center Working Paper, Institute of Business and Economic Research, University of California, Berkeley (2000).

Lewis, Matt, "Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market" Center for the Study of Energy Markets Working Paper, University of California Energy Institute, Berkeley (2003).

Lindley, D.V. and A.F.M. Smith, "Bayes Estimates for the Linear Model (With Discussion)" Journal of the Royal Statistical Society B 34 (1972), 1-41.

Manuszak, Mark D., "The Impact of Upstream Mergers on Retail Gasoline Markets" Working Paper, Graduate School of Industrial Administration, Carnegie Mellon University (December 2001).

Poirier, Dale J., "Intermediate Statistics and Econometrics: A Comparative Approach" The MIT Press, Cambridge (1995).

Rotemberg, Julio J. and Garth Saloner, "A Supergame-Theoretic Model of Price Wars during Booms" American Economic Review 76 (June 1986), 390-407.

Romley, John A., "Demand Complementarity in Retail Markets: Measuring the Effects of Convenience Goods on the Demand for Gas" Ph.D. Dissertation, Stanford University (December 2002).

Ross, Thomas W., "Cartel Stability and Product Differentiation" International Journal of Industrial Organization 10 (1992), 1-13.

Shepard, Andrea, "Price Discrimination and Retail Configuration" Journal of Political Economy 99 (February 1991), 30-53.

Slade, Margaret E., "Vancouver's Gasoline-Price Wars: An Empirical Exercise in Uncovering Supergame Strategies" Review of Economic Studies 59 (1992), 257-276.

Stigler, George J., "A Theory of Oligopoly" Journal of Political Economy 72 (1964), 44-61.

Tirole, Jean, "The Theory of Industrial Organization" The MIT Press, Cambridge (1988).