5-2017

Strategic Interaction in the Private and Public Sectors: Empirical and Theoretical Approaches

Peter D. Lorenz
Clemson University, plorenz88@gmail.com

Strategic Interaction in the Private and Public Sectors: Empirical and Theoretical Approaches

---

A Dissertation
Presented to
the Graduate School of
Clemson University

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy
Economics

---

by
Peter D Lorenz
May 2017

Accepted by:
Dr. Matthew S. Lewis, Committee Chair
Dr. F. Andrew Hanssen
Dr. Scott L. Baier
Dr. Raymond D. Sauer

Abstract

This thesis examines strategic interactions across multiple levels of society. From imperfectly competitive markets to political decisions, game theoretic results produce valuable predictions for economists. In the following three studies, I demonstrate some of the contributions and applications of modeling strategic behavior on the part of key decision makers. The first chapter provides new insights on the nature of competition and pass-through in retail gasoline markets. A large literature has identified tacit collusion and consumer search as explanations for the observation that gasoline prices rise more rapidly than they fall. This study contrasts the special case of geographically isolated stations against those in competitive markets. I find evidence that both tacit collusion and consumer search have an effect on pricing patterns, and that neither of these theories can entirely explain observed patterns of slow and asymmetric pass-through alone, as some have suggested. The second chapter investigates why company-operated gasoline stations adjust prices more quickly than franchised or licensed stations following a change to underlying costs. I present two theoretical explanations for this phenomenon. First, some firms may have less access to information about future market conditions and use competitors' prices as cost signals. Second, well-established corporate chains may take an active role in coordinating prices in otherwise competitive markets. The third and final chapter explores strategic interaction at the political level. The decision by a political body to expand the voting franchise can be seen as a strategic move to allow for a more favorable voting outcome in the future. Using a model that accounts for both voters' and policymakers' preferences, as well as some historical anecdotes, I study the case of an expansion of the voting franchise to younger voters. My results suggest that politicians who unexpectedly find themselves in a position of power will consider adding younger voters in order to maintain their position.

Table of Contents

List of Tables

Chapter 1

Price Fluctuations and Market Power: Evidence from Retail Gasoline Markets

1.1 Introduction

There exists a significant literature investigating competition and pricing patterns in the retail gasoline market. These studies examine many aspects of firm behavior: the pass-through of wholesale cost shocks onto retail prices, the asymmetric adjustment with which pass-through takes place, the influence of tacit collusion, and price dispersion across competitors. The inferences drawn from this literature have contributed to a broad understanding of competition in this market. However, competing explanatory theories remain, and important questions with regard to imperfect coordination and incomplete information have yet to be resolved. In part, this may be because the extent of empirical investigation has largely been limited to stations in fairly direct oligopoly competition, and in areas with many competing stations. In this paper I investigate station behavior within several particular and extreme market structure settings: single isolated stations and isolated pairs of competing stations.

One of the most studied pricing phenomena in retail gasoline markets is asymmetric price adjustment, whereby firms raise prices more quickly than they lower prices (Borenstein, Cameron, and Gilbert, 1997). Two primary explanations have been offered, one based on consumers' search behavior and the other based on tacit collusion between firms. The first of these theories, formalized by Tappata (2009), Yang and Ye (2008), and Lewis (2011), proposes that consumers engage in more searching when prices are rising than when prices are falling in order to learn about the nature of these price changes. The second theory, hypothesized by Borenstein et al., relies on a specific form of tacit collusion in which firms are unable to effectively coordinate when prices are increasing, but can use past prices as a "focal point "at which to collude when prices are falling. Stations adopt a strategy of not lowering their price until a rival changes its price, and are therefore able to temporarily maintain higher price margins after costs have fallen. Empirical evidence has been found in support of both explanations, but it remains unclear to what extent these mechanisms affect price pass through more generally. Moreover, while consumer search and tacit collusion provide reasonable explanations for asymmetric price adjustment in a typical competitive setting, they have different implications for stations with little to no nearby competition.

There are several predictions for how stations with no nearby competitors will behave based on the two primary explanations for price adjustment asymmetry. The consumer search explanation applies to isolated gasoline stations, provided they are not true monopolies. That is, as long as the market is covered and these isolated stations experience non-zero elasticity of demand, they face at least some threat of consumer search and will be slow to adjust price downward. However, the tacit collusion explanation does not apply to such stations; therefore, any asymmetry found in these stations would be best attributed to consumer search. The predictions for isolated pairs of stations are less straightforward. The consumer search must be applied differently depending on whether these stations are able to collude. In a collusive scenario, the firms should adjust price like a single isolated station. Thus, if the theory holds as an explanation for asymmetry, collusion should take place mainly when prices are falling, and isolated pairs should lower prices at the same rate as isolated single stations. However, the difference between isolated stations and isolated pairs when prices are rising is ambiguous. This ambiguity rests mainly on the trade-off between higher per-unit margins and market share. For a station with exactly one competitor who experiences an unexpected cost increase, a price increase keeps per-unit profit margins at the same level as previously, but it also induces customers to switch to the nearby competitor that may not have increased price.

I find that for both price increases and decreases, competitive stations adjust prices more quickly than do geographically isolated stations with more market power, something that is not found in the price adjustment literature to this point. This result does not necessarily contradict either the tacit collusion or the consumer search models for price adjustment asymmetry, although it does require some additional explanation for isolated stations' slow price increases. I find that isolated pairs of stations are even slower than isolated single stations to increase price. This lends credibility to the tacit collusion model, which suggests that nearby stations use current prices as a focal point for collusive behavior. As such, collusion is easier when prices are falling than when they are rising. My result suggests that on average, the cost to raising price above a single competitor exceeds the cost of keeping prices below the long-run cooperative price. However, these isolated pairs reduce price at about the same speed as isolated single stations, suggesting that when costs are falling, they are able to collude on price and act like a single isolated station.

A second contribution of this paper is the introduction of a more appropriate method of empirically modeling station price adjustment. Existing studies have almost exclusively utilized a vector-autoregressive or error-correction approach to modeling cost pass-through. Unfortunately, these models impose a gradual adjustment through incremental changes in each period. While this may serve as an effective approximation when studying aggregate data, it does not accurately reflect the way in which individual stations implement price changes. On average, stations in my sample only change price on about one in every five days, choosing to keep price unchanged otherwise. For this reason, I separately model station's likelihood of a price change from the magnitude of a price change.

The rest of the paper will proceed as follows. In Section 2, I review various explanations for price rigidity and asymmetry, and explain the motivation behind my analysis. In Section 3, I present a theoretical explanation for slow price pass through, providing separate explanations for slow price increases and for slow price decreases. In Section 4, I describe the data used for this study and discuss the econometric model. In Section 5, I report the results of the empirical analysis. In Section 6, I conclude the study.

1.2 Motivation and Previous Studies

It is important to recognize the findings of previous works involving price rigidity, market power in the gasoline market, and price adjustment asymmetry. In one instance of an analysis focused on market power, Hong and Lee (2014) approached the phenomenon of price adjustment asymmetry using geographically separated gasoline stations on Korean islands. They found that market power increased the degree of asymmetric adjustment. In addition, they tested various explanations for this observed asymmetry. Their evidence supported both the consumer search and tacit collusion explanations. Byrne (2015) has also addressed the problem of pricing asymmetry in rural and urban markets. Using retail data from Canadian gas stations, that paper finds increased asymmetry in rural markets in which collusion would be more attainable. Both of these papers highlight the increased disparity between price response upward and downward. However, these papers do little to address the sluggishness of price changes for both increases *and* decreases in response to underlying costs among stations with few competitors. In addition, they do not apply the tacit collusion and consumer search models directly to isolated markets. This paper analyzes each of these individual cases separately in order to achieve a more well-rounded understanding of retail price dynamics.

Davis and Hamilton (2003) studied price stickiness in the retail gasoline markets using a model based on Dixit's (1991) theory of administrative costs. Davis and Hamilton applied this model, which used ongoing uncertainty and costly reversibility, to explain price stickiness to nine gas wholesalers in Philadelphia. They found that stickiness was likely due to firm expectations about consumer and competitor reactions; they rejected a literal translation of the menu cost model of administrative costs, favoring a broader interpretation of costs associated with price changes. Under this assertion, we would only expect to see slower-than-normal pass-through in isolated stations if these stations were especially worried about responses from loyal customers. This result may apply more to isolated pairs of stations who may attempt to reach a collusive price in a coordination game.

Bresnahan and Reiss (1991) discussed the role of entry into oligopolistic markets. Using a sample of isolated local markets, they examined the competitive effects of adding one firm to a small cluster of competing firms. They found that - contrary to expectations - most industries revert to competitive pricing after surpassing a low threshold of firms, as opposed to a gradual shift from monopolistic behavior to perfectly competitive behavior. This result suggests that price pass-through in gasoline markets may differ in the case of a monopolist from the case in which there are three to five firms, but it also predicts that a much larger number of firms may not be distinguishable from such smaller numbers. For this reason, I classify stations with at least five nearby competitors to be competitive.

1.3 Theory

The first group of firms I consider is gasoline stations that are geographically isolated from competing stations. Specifically, I refer to any station with no competitors within five miles (as the crow flies) as "isolated."In this way, I identify stations with a high level of market power. However, these stations are not true monopolies. While they have a very low cross-elasticity of demand with respect to competitors five or more miles away, it is also the case that there will be marginal consumers between an isolated station and its nearest competitor. For this reason, I expect nearly full price pass through in the long run for even the most isolated stations.

Similarly, isolated pairs of stations are not true duopolies, though they possess some duopolistic characteristics. Stations that are part of an isolated pair have a high cross-price elasticity with one another, but a low cross-price elasticity with the next closest station, which by definition is at least five miles away. Using the same logic as above, an isolated pair engaging in tacit collusion will be, at best, behaving like an isolated station, and *not* like a monopoly. That is, the pair's optimal strategy in perfect coordination is to set a price somewhat lower than the monopoly price, but higher than a competitive price.

Stations with many nearby competitors, henceforth referred to as "crowded"stations, are likely to have high cross-price elasticities with multiple competitors, since consumers experience low search costs. Thus, I expect these stations to behave in the way that any competitive firm would. Their long-run price should be a competitive price close to the underlying per-gallon cost.

According to the previous literature on price response asymmetry, prices rise after a cost increase more rapidly than they decline after a cost decrease (Bacon, 1991; Karrenbrock, 1991). The consumer search explanation for this phenomenon has implications on how market power affects price pass-through. The reasoning is as follows[1]: When costs are rising, firms are pressured to raise price in order to avoid negative profits. Some consumers respond to rising prices by searching for lower priced competitors, but they find that the only alternative is another higher-priced station. When costs fall, stations should be similarly pressured by competitive forces to lower their prices. However, if information about competing stations is scarce, then consumers don't search for lower prices; they don't know that lower prices could exist, because they don't realize that underlying costs have fallen. When a station gradually lowers its price over time, its regular consumers treat this as

[1] (Yang and Ye, 2008; Tappata, 2009; Lewis, 2011)

an unanticipated benefit, and they again are unlikely to search for lower prices. Because consumer search leads to asymmetric competitive pressure, prices rise in response to cost shifts more quickly than they decline.

Tacit collusion has also been considered (Karrenbrock, 1991; Borenstein et al., 1997) a possible explanation for asymmetric price adjustment: If there is a sufficiently small number of stations competing in a market, those firms can coordinate using "focal point" pricing, in which stations tend toward particular prices such as the most recent price posted. When costs are falling, these stations can keep prices higher than the competitive price, splitting total profits close to what a monopoly would earn. When costs rise, competing stations find coordination difficult. If tacitly collusive stations wish to raise price simultaneously in response to a cost increase, it may be difficult for them to know exactly what price they should charge, since past prices are not a useful focal point when costs are rising. As a result, the competing stations forgo collusion and raise price quickly. When costs fall, these stations drop prices slowly together.

I integrate both of these theories into my explanation for variation in price rigidity. Consumer search drives the difference between cross-price elasticities when prices are rising and when they are falling. Tacit collusion explains why some stations might be hesitant to change price at all, and why competition will not necessarily drive prices downward when costs are falling. If the losses from having a higher price than one's competitors outweigh the losses from having a smaller profit margin, the firm will refrain from increasing price in the presence of rising costs. If, on the other hand, the losses from smaller margins outweigh the losses from having a higher price than one's competitors, the firm will raise price relatively quickly. In addition, the transition from a collusive price outcome to a competitive environment could take time.

In the case of isolated stations, the tacit collusion model can be safely ignored. Only the consumer search explanation is in play, and the asymmetry found in these stations serves to support this assertion; without some kind of competitive effect from far away stations, there is little explanation for asymmetric price adjustment among stations with no nearby competitors. When costs rise for isolated stations, simple price theory indicates that they will adjust price upward fairly quickly. When costs fall for isolated stations, these stations should lower their price quickly *only* if they face zero competition. Since isolated stations are *not* true monopolies, they operate in equilibrium at a price somewhat below the monopoly price. Thus, when costs are falling, they can take advantage of a temporary lack of consumer search by keeping price higher than equilibrium, and closer to the true monopoly price.

For isolated pairs of stations, the tacit collusion model can be applied. If they are colluding, we would expect that they behave similarly to isolated single stations while prices are falling. If in the long run, pairs are able to set a collusive price equal to that of an isolated station, then they should respond to a cost decrease by holding that price steady while consumers fail to search for a cheaper alternative. However, when costs increase, these pairs may respond slowly if immediate coordination is difficult. In measuring general price rigidity, I am able to test how quickly isolated station pairs adjust from a tacitly collusive condition to a competitive oligopoly.

For crowded stations, a tacit collusion model again is not applicable, since there are likely too many firms to coordinate. However, I would expect that consumer search would have a different impact on these stations than it does on isolated stations. For isolated stations, consumer search is extremely costly, causing them to change price more slowly in each direction. Crowded stations experience low-cost consumer search and possibly lower profit margins due to increased competition. When costs increase, these stations must increase price to cover cost. When costs fall, I expect crowded stations to reduce price more quickly than isolated stations due to the low cost of price comparison for consumers. However, I also expect prices to fall more slowly than they rise at crowded stations in accordance with previous findings regarding price adjustment asymmetry.

1.4 Data and Empirical Approach

I used daily retail gasoline price data gathered from the Oil Price Information Service, a leading price information provider within the industry. Data were collected from six different states: New York, Pennsylvania, Kansas, Oklahoma, Alabama, and Arkansas. These states were chosen due to their proximity to the oil import harbors (the Gulf Coast and the New York Harbor), as well as their mix of rural, urban, and suburban areas. In this way, I acquire data on very small towns with one or two gas stations, as well as areas with many more stations. The unit of observation is the station-day. Each observation contains the station's name and postal address, a unique station ID number, the station's latitude and longitude, the date and time of collection, and the price posted at the time of collection.

In addition, I acquired cost data from the U.S. Energy Information Administration. I used the New York Harbor conventional gasoline regular spot price as a proxy for the underlying cost of gasoline for stations in Pennsylvania. I used the U.S. Gulf Coast spot price for stations in the other four selected states.

I used the longitude and latitude information to identify the number of neighbors each station has within a particular pre-determined distance. I applied Vincenty's formula [2] to calculate the distance from each station to each other station. I then define isolated stations as those with no competitors within five miles. I define an isolated pair in the following way: Station j is a part of an isolated pair if has exactly one neighbor within 5 miles, that neighbor is also within 1 mile of station j, *and* that neighbor's only neighbor within 5 miles is station j. In this way, I identify pairs of stations who are within 1 mile of one another, but not within 5 miles of any other station. This identification should allow me to measure the marginal effect of adding a single competitor to an otherwise monopolistic scenario. As defined above, a crowded station is any station with at least five competitors within 1 mile.

In a preliminary analysis, I find that during my sample period, isolated stations and isolated pairs changed price significantly less frequently than crowded stations, in the cases of price increases and decreases. Tables 1 and 2 describe the frequency of price changes for gasoline stations of various competition types in the sample. An interesting observation is the low frequency with which *all* stations change price. The typical gas station changed price on only about one in five days. Note that price decrease frequencies are larger than price increase frequencies. This does not contradict the known price asymmetry effect, because costs were falling more often than they were rising during my sample period. However, comparing across levels of competition can still be instructive.

Tables 3 and 4 describe the average magnitude of price increases and decreases for each market, given a price change $\neq 0$. While crowded stations appear to increase prices by larger amounts and decrease prices by smaller amounts than less competitive markets, it is difficult to make any conclusions about firm behavior from these statistics. In order to make such conclusions, I must conduct a thorough analysis of stations' responses to individual observed cost shocks.

While my measure of cost does not include transport costs or other non-fuel marginal costs, these are likely to be time-invariant and will be absorbed by station fixed effects. Station fixed effects will also capture any persistent differences in profit margins across stations that will arise from differences in the degree of local competition. Following Lewis (2011) and others modeling

[2] This method, developed by Vincenty (1975) calculates the distance between two points on a sphere's surface, and is useful in calculating distances over the Earth's surface.

gasoline price adjustment, I utilize an error correction model and a two-step estimation approach similar to Engle and Granger (1987). In the first step, I estimate the long-run relationship between price and cost using a station fixed effects model.

$$p_{jt} = \alpha_j + \phi c_{jt} + \mu_{jt} \tag{1.1}$$

I use the residual from the first stage regression as an estimate for a station's deviation from its expected markup. I then model a station's daily price change as a function of that deviation and of any recent cost changes.

$$\Delta p_{jt} = \sum_{s=1}^{10} \beta_s \Delta c_{t-s} + \theta(\hat{p}_{jt} - \alpha_j - \phi c_{jt}) \tag{1.2}$$

or,

$$\Delta p_{jt} = \sum_{s=1}^{10} \beta_s \Delta c_{t-s} + \theta \hat{\mu}_{jt} \tag{1.3}$$

Δc_{t-s} represents a set of lagged cost response term for s days prior to day t. This is to account for any tendency to change prices sooner rather than later. These effects statistically drop to zero after about 10 days in all cases, so I have restricted the lagged cost responses to 10 days.

When μ_{jt} is positive, its magnitude will strongly impact the amount by which stations lower their price, as well as their decision to do so. Similarly, when μ_{jt} is negative, its magnitude will strongly influence the size and frequency of price increases. However, it is likely the case that when μ_{jt} is positive, its magnitude will not strongly affect the size of a price increase, since such increases are infrequent and incidental. For the same reason, when μ_{jt} is negative, its magnitude may not strongly impact the size of a price decrease. To capture this, I employ a threshold autoregressive model (Enders and Granger, 1998) in which I create a dummy for whether a station's margin is above its expected margin at that cost:

$$I = \begin{cases} 1 & \text{if } \mu_{jt} < 0 \\ 0 & \text{otherwise} \end{cases} \tag{1.4}$$

Further, I must separately estimate a station's likelihood of changing price at all from the magnitude of its price change. Following this approach, I develop a logistic probability model for the binary decision to change price. We have:

$$\Psi_{jt}^{I} = \gamma_0^I + \sum_{s=1}^{10} \gamma_s^I \Delta c_{t-s} + \tau^I I + \lambda^{IL} I \mu_{jt} + \lambda^{IH}(1-I)\mu_{jt} + \epsilon_{jt}^I \tag{1.5}$$

for price increases, and

$$\Psi_{jt}^{D} = \gamma_0^D + \sum_{s=1}^{10} \gamma_s^D \Delta c_{t-s} + \tau^D I + \lambda^{DL} I \mu_{jt} + \lambda^{DH}(1-I)\mu_{jt} + \epsilon_{jt}^D \tag{1.6}$$

for price decreases, where Ψ refers to the latent measure of the incentive to change price in a logit model, θ^L refers to the firm's price response when $I = 0$, and θ^H refers to its price response when $I = 1$. To estimate the magnitude of these price changes *given* the decision to change price at all, I have

$$\Delta p_{jt}^{I} = \beta_0^I + \sum_{s=1}^{10} \beta_s^I \Delta c_{t-s} + \nu^I I + \theta^{IL} I \mu_{jt} + \theta^{IH}(1-I)\mu_{jt} + \varepsilon_{jt}^I \tag{1.7}$$

for price increases, and

$$\Delta p_{jt}^{D} = \beta_0^D + \sum_{s=1}^{10} \beta_s^D \Delta c_{t-s} + \nu^D I + \theta^{DL} I \mu_{jt} + \theta^{DH}(1-I)\mu_{jt} + \varepsilon_{jt}^D \tag{1.8}$$

for price decreases. The above equations were estimated conditional on the presence of a price change. That is, (7) and (8) assume that $\Delta p_{jt} \neq 0$.

1.5 Results

Using the 2-step estimation technique described above, I record the relevant parameter estimates in Table 5, separated by market type. Note that the differences between isolated and crowded stations' responses to cost changes appear very similar. The primary driver for the difference in their rate of price change can be found in the lagged cost response terms, $[\gamma_1, ..., \gamma_{10}]$, which can be found in Appendix A. I do find that the difference in θ_L for isolated and crowded stations can help explain price pass-through differences when prices are increasing. For isolated pairs, the differences in response to cost shocks are more stark. In response to cost increases and decreases, these stations are less likely to change price for any given cost shock. The price change magnitude response for isolated pairs is slightly smaller for price decreases than other market types. For price increases, it's nearly the same as for isolated stations.

In order to properly analyze how these parameters affect price pass-through, it is necessary to simulate long-run price responses to a cost shock. I use a series of 10,000 bootstrapped simulations to accomplish this. Each simulation is assigned a set of parameters $[\theta, \lambda, \beta_0, ..., \beta_{10}, \gamma_0, ..., \gamma_{10}]$ randomly chosen from the distribution of my coefficient estimates. Within each simulation, for each market type, I impose a sudden cost increase of 25 cents. Since my parameter estimates were made with respect to a station's predicted price $\hat{p}_{jt}$, I can normalize each station's initial markup as $\mu_{j1} = 0$. I then simulate a 30-day follow-up period with 10,000 stations. For each day, each station's current markup μ_{jt} is measured. The station then chooses to increase price, or not, with probability Pr(Inc) according to equation (7). If the station chooses to increase price, it does so by Δp^I_{jt} as described in equation (5). For each day, I record the average price of the 10,000 stations in that iteration of the simulation. The result is a series of 10,000 averages over 30 days, with each average representing a particular randomly generated set of parameters. I then record the average of these averages, as well as the 5th and 95th percentile of these averages. I conduct a similar process for price decreases.

Figure 1 shows the estimated price paths for isolated stations, isolated pairs, and crowded stations after a cost increase, as well as 95 percent confidence intervals for each. Crowded stations experience the fastest upward adjustment. They reach 80 percent pass-through by day 8 after the cost shock. Isolated stations reach 80 percent pass-through on day 15, while it takes 24 days for isolated pairs to reach that mark. This contrast demonstrates that isolated stations are slower to respond to cost increases than crowded stations. As outlined above, there are few explanations for this outside of a menu cost or information cost explanation. Further investigation of the causes of this particular phenomenon is left to future research.

However, the low price adjustment exhibited by isolated pairs suggests that these stations struggle to coordinate in the short run on price increases. Because they raise prices even more slowly than isolated stations, I posit that the losses these stations would incur from having a higher price than the nearest competitor outweighs the gains from an increased profit margin. In the long run, these stations do pass through costs to their prices, eventually reaching a coordinated equilibrium.

Figure 1.1

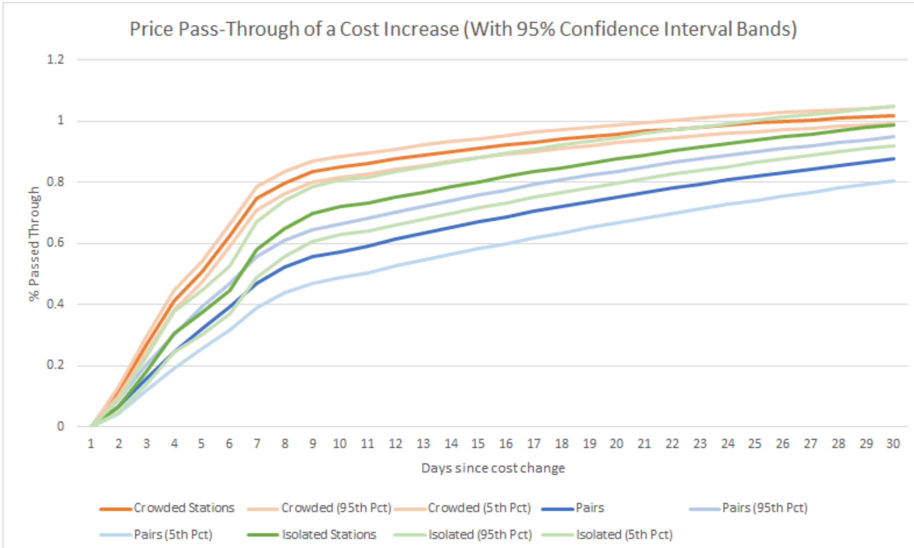

Figure 2 shows estimated price paths after a price decrease. Crowded stations again exhibit the speediest price adjustment, reaching 80 percent pass-through by day 18. Note that this is considerably slower than the upward pass-through for the same market type. This is consistent with previous work on price adjustment asymmetry. Isolated stations and isolated pairs are both slower to adjust than crowded stations, reaching 80 percent pass-through on day 24, but their price paths are nearly identical to one another. This suggests that station pairs are able to collude more easily when prices are falling, exhibiting the same price path as a single station with a high degree of market power.

Figure 1.2

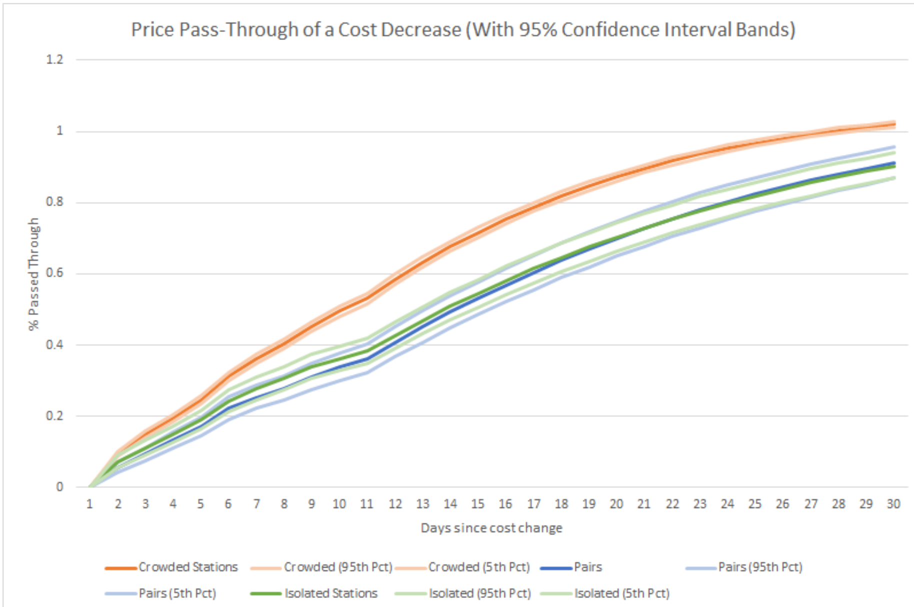

This comparison between isolated stations and station pairs supports the claim that retail gasoline stations use focal point prices to engage in tacit collusion. Specifically, it supports the claim that they use the most recently posted price as a point of collusion. When prices are falling, these stations may find coordination easier by simply remaining at a previously posted price. Consumer search is relatively low during a cost decrease, because consumers are unaware that market conditions have changed. They observe unchanged prices at the only two sellers nearby. Gradually, these stations drop price and eventually pass prices through.

1.6 Conclusion

Earlier work in the study of menu costs and gasoline markets have indicated the possibility of a relationship between price stickiness and market power, but have never made a definitive statement about the direction or magnitude of this relationship. In addition, very little work has been done on geographically isolated stations with very little competition. My results show that such isolated stations are slower than more competitive markets to raise and lower prices in response to exogenous cost shocks.

Furthermore, I show that near-duopolistic stations with no nearby competitors except one another raise prices even more slowly than single stations in the same geographic setting, despite lowering prices at the same (very slow) rate. These results support earlier explanations for price response asymmetry: low consumer search during cost decreases, and tacit collusion. My results *also* indicate that collusive behavior is difficult to conduct when costs are rising. Instead of immediately raising price to the optimal near-monopolistic price that isolated stations reach, these station pairs remain at a lower price for a brief transition period from a collusive setting to an oligopolistic one.

My results indicate that both the consumer search and tacit collusion explanations for asymmetric price adjustment are valid. I observe evidence for each of these phenomena separately, whereas previous studies attempted to test between these two and favor one over the other. However, neither of those theories can single-handedly explain the above results.

Chapter 2

Company Operated Firms and Competitive Effects in Retail Gasoline Markets

2.1 Introduction

In addition to the literature investigating competition and pricing patterns in the retail gasoline market, there has also been a large discussion on the effects of franchised firms and vertical integration. Much of this work has been focused on the principle-agent problem and moral hazard surrounding the employment of local managers who are charged with operating and maintaining a company store, as well as the positive effects that franchising has on company growth and expansion. However, there has emerged a smaller literature on the impact of franchised firms on local competitors' pricing behavior. Under the assumption that major corporate entities possess better information regarding future market conditions, it is likely that company-owned firms adjust more quickly to upstream cost shocks than independent or franchised firms. If those independent competitors are aware of this information asymmetry, they may choose to take franchised competitors' prices as a signal of market conditions that are otherwise difficult to observe. This paper will examine the possibility that competing gasoline stations engage in a leader-follower pattern, in which the most well-informed station acts first in response to cost shocks about which its competitors may not be aware, and in response, those competitors adjust their retail prices to account for the new information observed through the well-informed station.

Early work on franchisee behavior has shown that such firms set prices and promotions in a somewhat formulaic manner in response to upstream parent companies. Moore and Sherwood (1963) describe the effects of widespread franchising in the soft-serve ice cream industry, showing that corporate influence allowed for more growth and uniformity among chained firms: "sales promotion policy also appeared to be affected by the franchise. Soft-serve firms followed promotion policy of product differentiation. However, firms franchised by the same organization used similar if not identical advertising and point-of-sale promotions. This action altered behavior from that expected had each firm acted independently."In this context, the franchise was contrasted to independent firms with no brand affiliation at all. In other words, increased corporate involvement in firm-level decisions resulted in distinct and predictable differences in firm behavior.

This paper extends this concept to company-operated firms, which should exhibit even more uniformity than a franchised firm. I hypothesize that a corporate structure that can acquire information about market conditions may be able to direct downstream firms in a way that, for the purposes of this study, allows those company-operated firms to behave generally like well-informed firms. Because such firms act uniformly and quickly, competing firms may adjust their approach to setting prices such that the franchised firm's price becomes something of a "focal point"for other firms' pricing. That is, those competing firms could use a well-informed neighbor's posted price as a signal for upcoming market conditions, as opposed to obtaining that information themselves at a relatively high cost. In order to determine whether this activity is taking place, I analyze the retail price responses of a large sample of gasoline stations in the United States under relatively high levels of competition. categorized as one of three possible station types: 1) A company-operated QuikTrip, Casey's General Store, or Sheetz gas station; 2) A nearby competitor of a QuikTrip, Casey's General Store, or Sheetz location; or 3) A station that has nearby competitors, but none of which is a QuikTrip, Casey's General Store, or Sheetz. These three companies are among the largest chains of exclusively company-operated locations.

In order to study these firms' behavior, I must account for known asymmetry in price responses. In retail gasoline pricing, it has been found that prices increase more quickly than they decrease in response to changes in underlying costs. Additionally, I find that price changes take place relatively infrequently; on average, a gasoline retailer changes ts price once every five days. For these reasons, I separate price increases from price decreases entirely, treating them as independent choices. This is in contrast to other price adjustment models, which treat price as a continuously updated choice variable using an error correction approach to empirics. Instead, I treat a price adjustment as a discrete daily decision by the firm before determining the magnitude of such a change.

I find that selected company-operated firms tend to increase prices very quickly in response to observed cost shocks, relative to stations with similar levels of competition. QuikTrip and Sheetz also tend to decrease prices more quickly than comparable firms, while Casey's stations decrease price at an unremarkable rate. In both cases, QuikTrip stations adjust prices the fastest. Furthermore, nearby competitors of QuikTrip locations tend to increase prices at a rate just below QuikTrip stations, indicating a leader-follower dynamic consistent with my information asymmetry hypothesis. Neighbors of Sheetz and Casey's locations also follow price paths similar to their respective company-operated competitors. QuikTrip's neighbors also decrease prices at a higher rate than the typical competitive firm, although this effect is less pronounced than it is for price increases. Sheetz competitors exhibit somewhat faster pass-through than the typical competitive firm. Casey's neighbors pass through cost decreases at a rate similar to the typical competitive firm, which is consistent with Casey's stations behaving in that way when prices are falling.

These results provide evidence that major company-operated retailers act as price adjustment leaders in local gasoline markets, leading to very fast price adjustments. I consider two explanations for observed price leadership. First, company-operated firms may be better informed about future supply conditions, resulting in faster price response to cost shocks. These price adjustments provide nearby gasoline stations with valuable information about upcoming changes in the underlying costs of gasoline. Such competitors respond to a change in a company-operated retailer's posted price by matching that price change. Although many of these firms do not have access to the information needed to make such an adjustment (or choose not to acquire it at a cost), that information is spread through price changes among well-informed retailers. A second explanation considers a scenario in which all firms are well-informed, but they are unable to coordinate price changes shortly after a cost change. Company-operated stations in this case may take on the role of a price coordinator in a competitive market.

The rest of the paper will proceed as follows. In Section 2, I review the literature on gasoline pricing dynamics and company-operated firms, and explain the motivation behind my analysis. In Section 3, I present a theoretical model of information asymmetry and leader-follower interaction.

In Section 4, I describe the data used for this study and discuss the econometric model. In Section 5, I report the results of the empirical analysis. In Section 6, I conclude the study.

2.2 Motivation and Previous Studies

Due to product homogeneity and the relatively low cost of entry, retail gasoline stations have long been considered a strong example of sellers in a competitive market. Specifically, they are often assumed to engage in Bertrand-style competition, reducing economic profits to near zero. However, the degree to which gasoline retail markets exhibit non-competitive or oligopolistic tendencies has been studied carefully for nearly three decades. Shepard (1990) discussed a variety of explanations for observed high margins, including tacit collusion, price discrimination, and double marginalization. That paper also investigated the nature of contracts between manufacturers and retailers designed to minimize principle-agent problems. In 1991, Shepard expanded upon the price discrimination hypothesis and found evidence that quality-based price discrimination added at least 9 cents to the price of a gallon of gasoline.

Hastings(2004) found that profit margins for stations selling branded gasoline was directly related to the brand itself, rather than the vertical structure of the stations. That is, company-operated branded stations had the same impact on local prices as did a franchised or dealer-owned station selling branded gasoline. However, she did find evidence that the introduction of independent non-branded stations drove local prices downward, since such stations compete only on price. The company-operated chains in my study have similarities to each of these types of stations. Like Hastings' independent stations, these chains sell unbranded gasoline and compete mainly on price. However, like Hastings' company-operated branded stations, these chains also offer a large number of indoor amenities and products. This provides an incentive for stations like QuikTrip, Casey's General Store, and Sheetz to offer low-priced unbranded gasoline in order to increase their volume of customers and raise indoor sales.

Hastings' result implies that in a leader-follower scenario, stations like QuikTrip, Sheetz, and Casey's are likely to decrease prices the fastest in response to a cost decrease, since they are primarily concerned with increasing customer volume. In the case of a cost increase, it is less clear that a company-operated station would be the first to increase price. If the station believes its competitors will follow suit, then an increase in price will strictly increase margins on gasoline, and will not have a strong effect on indoor traffic as long as the gasoline price relative to other local prices does not change. If the station believes a price increase will not be met by other competitors, then it loses customers for both gasoline and indoor amenities in classical Bertrand fashion.

Noel (2007) discussed the role of smaller gasoline stations in areas with Edgeworth cycles, in which prices rise very sharply, followed by a gradual decline in prices over the course of several weeks. He found that cycles and sticky pricing are more prevalent in areas with many smaller firms. The literature on Edgeworth cycles, which started with Maskin and Tirole (1988) and has received further contributions from Doyle et al. (2009) and Lewis and Noel (2011), suggests that smaller independent firms in Edgeworth Cycles are following closely behind stations that are part of larger corporations, especially when certain companies have several locations within a cycling area. This behavior could be caused by imperfect information among smaller firms who use larger firms' prices as reference points for optimal pricing.

This paper attempts to identify the order in which competing company-operated and non-company-operated stations change price following an observable cost change, as well as the speed at which other competitive stations respond to such a change. The primary contribution of this paper is to determine whether these differences are due to firm preferences or asymmetric information, and in the case of the latter, if well-informed firms' prices are used as cost signals by less-informed firms.

2.3 Theory

Because this paper is concerned primarily with how the existence of a company-operated firm affects local prices, I ignore stations with fewer than 5 competitors within one mile as the crow flies. After conditioning my sample on stations with such a level of competition, I separate gasoline stations into three groups. The first group of firms is the set of locations owned and operated by QuikTrip, Casey's General Store, and Sheetz. These three companies (as well as some other common chains that do not appear frequently within the geographic confines of my data set) often have price determined upstream at the corporate level. The second group of stations are stations located within one mile of a QuikTrip, Casey's General Store, or Sheetz location, but which are not one of those chains themselves. These stations are referred to as QuikTrip Neighbors, Casey's Neighbors, and Sheetz Neighbors respectively. The third group of stations are those which meet the requisite level of competition, but which are neither near one of the three specified company-operated chains, nor are they one of those company-operated chains themselves.

I consider two distinct models for gasoline markets to describe the potential price leadership of company-operated stations on local competition. The first involves incomplete information on the part of stations that are not operated by a large parent company. The second treats company-operated firms' prices as focal points in larger, more competitive markets in which price coordination is otherwise difficult.

While a significant literature has discussed the effects of asymmetric information between buyers and sellers,[1] there has been less work describing non-uniform information among firms. Because information about future supply and demand conditions is costly, not all stations choose to acquire such information. As a result some stations may lag behind others in responding to changing market conditions. However, uninformed firms that can observe the price set by well-informed firms can take that as a signal for changing market conditions and adjust their own prices accordingly.

In an information asymmetry model, I consider company-operated stations to be well-informed about future cost conditions. I consider all other stations to be less informed. This model has specific predictions about information flow in markets with a well informed station in competition with less informed stations. For example, in a competitive market that features one QuikTrip location and several non-company-operated stations, the QuikTrip station acquires information frequently at a low cost, and it adjusts its price accordingly. QuikTrip's neighbors acquire information far less frequently due to high costs of acquisition. However, they do observe a price change at the local QuikTrip. If these stations believe that QuikTrip will reliably make optimal price adjustments, their own information costs are close to zero after observing QuikTrip's price. They will therefore change price almost as frequently as the QuikTrip, but they will not do so until after the QuikTrip has already changed price.

Independent or franchised stations in a competitive setting with no company-operated neighbors will behave differently. With a high cost of information acquisition and no well-informed stations nearby to provide the relevant information, these stations are likely to let prices remain unchanged for a longer period of time. I would expect these firms to pass through cost shocks at a considerably lower rate than the other two types of station.

A key component of these theoretical results is that while price response asymmetry (prices

[1] The classic works of Akerlof (1970) and Arrow (1963) describe situations in which either the buyer or seller has more information than the other about the value of a good being transacted.

rise faster than they fall) could still take place, the differences between informed stations, their neighbors, and non-neighbors should persist across price increases and price decreases. If empirical analysis only shows company-operated stations leading the pack when prices are falling (and not when they are rising), then although they may be taking advantage of an information advantage, their neighbors are simply responding to Bertrand market pressure to lower prices. If company-operated firms respond to cost shocks more quickly than other firms, but their neighbors respond no more quickly than other competitive stations, then the above theory is not supported; there are many reasons that certain types of firms would pass cost shocks onto prices more quickly than others.

Another explanation for price leadership by corporate chains is that these company-operated stations act as price coordinators in a market of perfectly informed firms. Using this approach, the kind of tacit collusion often seen in smaller markets[2] is difficult to achieve with many firms in a small geographic area. Because companies like QuikTrip, Sheetz, and Casey's General Store often post lower gasoline prices than nearby competitors, a modest price increase might still result in a relatively low price. As for price decreases, stations that are primarily concerned with having a lower price than their competitors (including but not limited to these company-operated firms) may have a strong incentive to respond promptly to a cost decrease.

These two theories could both result in price leadership by stations with large parent companies. However, they differ in some key implications. The asymmetric information explanation implies that a particular company-operated station should increase and decrease prices earlier than their neighbors by about the same amount of time, since the lag is caused by indpendent stationss failure to observe a well-informed station's price change in real time. That is, if it takes two days for a franchised Shell station to observe and respond to a local QuikTrip's price increase, it should take two days for that Shell station to observe and respond to the QuikTrip's price decrease. The price coordination explanation allows for that Shell station to respond more quickly to QuikTrip's price decrease than to its price increase, since there is no temporary gain from failing to coordinate when prices are falling. However, a coordination explanation does not entirely account for increased price pass-through in markets with a company-operated firm. While I would expect such markets to increase prices more quickly with stronger coordination, I would expect them to decrease prices at a similar rate to other competitive gasoline markets, since those markets do not require a specific

[2] This is discussed at greater length in the previous chapter

price-coordinating firm to drive down prices when costs are falling.

2.4 Empirical Approach and Results

Much of the Empirical approach to this study is carried out in the same manner as in the first chapter. I use data collected from various geographic locations somewhat near the Gulf Coast and New York Harbor. Unlike the previous study, I condition here on competition level. Only stations with at least 5 competitors within 1 mile are considered, so that all stations analyzed are in very competitive markets[3].

Because these competitive stations still exhibit fewer price changes than classical Bertrand competition would predict, I again consider price changes to be discrete day-to-day decisions that can be described in an error correction model by the equation:

$$\Delta p_{jt} = \sum_{s=1}^{10} \beta_s \Delta c_{t-s} + \theta \hat{\mu}_{jt} \tag{2.1}$$

In addition, because of known price response asymmetry, equations 5-8 from the previous chapter also apply; I estimate the probability of a price increase separately from the probability of a price decrease. In this study, I estimate those four equations for company-operated firms, their neighbors, and independent or franchised stations with no company-operated neighbors. I then use a series of simulations similar to the ones used in the first chapter.

My findings support my hypothesis that not only do major company-operated stations adjust prices more quickly, but that they provide valuable information about costs to smaller nearby firms with less access to that information. Figure 1 describes the price response path following a simulation of a 25-cent increase to gasoline harbor prices. QuikTrip and its nearby competitors adjust prices upward considerably faster than do other types of station, with QuikTrip leading its neighbors by a few days, and ultimately making a larger adjustment overall on average. Sheetz locations and their neighbors pass costs through more slowly than markets with a QuikTrip, but much more quickly than typical competitive stations. Casey's and its neighbors adjust slightly more slowly than Sheetz, but again more quckly than typical stations.

Figure 2 shows the price response path following a simulation of a 25-cent decrease to harbor

[3]As in the previous chapter, Bresnahan and Reiss point out that in markets with 5 or more firms, the marginal effect of an additional firm has a minimal effect on pricing.

Figure 2.1

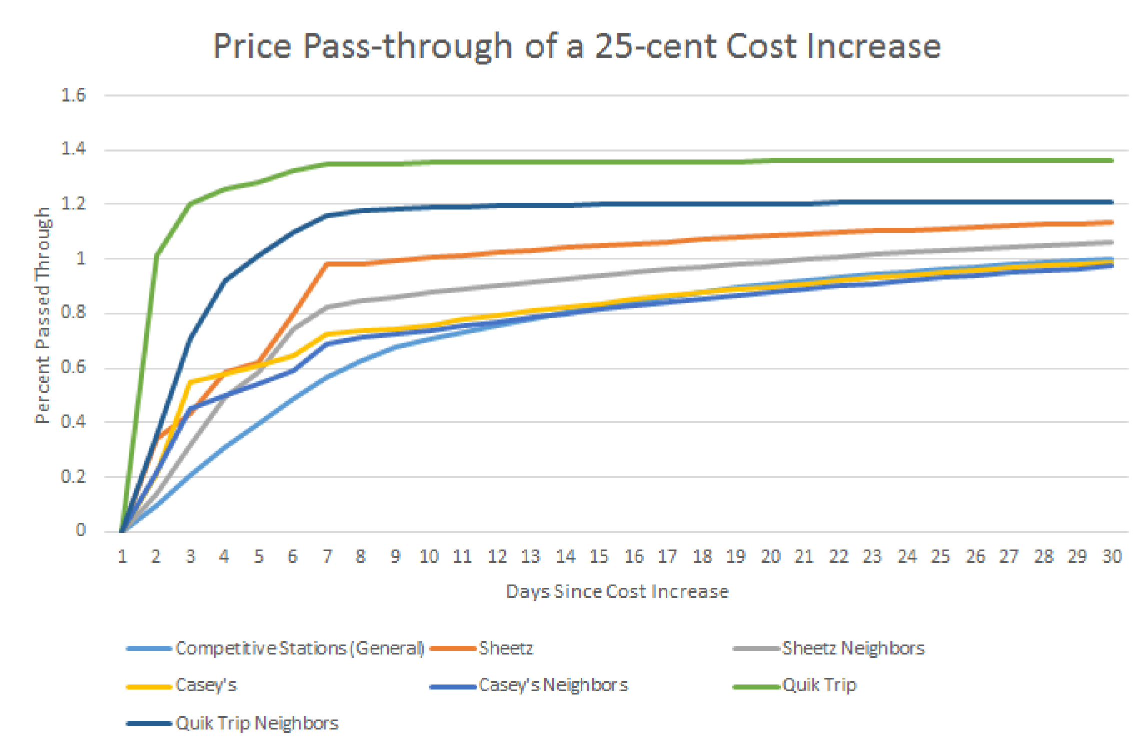

prices. Again, QuikTrip and its neighbors are the fastest to adjust. In this case, QuikTrip's neighbors follow much more closely to QuikTrip than they do for price increases. Sheetz stations are the next fastest company-operated chain to adjust to cost decreases, with their neighbors also following very closely. Casey's locations do not decrease prices at a rate that is significantly different from typical competitive stations, which allows their nearby competitors to behave in the same way.

These results provide evidence of both explanations for price leadership on the part of company-operated firms. The increased speed at which markets featuring a QuikTrip or Sheetz decrease price indicates that those firms are acting as information providers to their competitors. However, because their competitors do not lag behind by very much time, there is a strong case that price leadership exhibited in Figure 1 when prices are rising can be explained by imperfect coordination. When a QuikTrip location increases its price, its competitors are informed about an upcoming cost increase, but are hesitant to increase price immediately due to the presence of other stations that have yet to increase price.

Figure 2.2

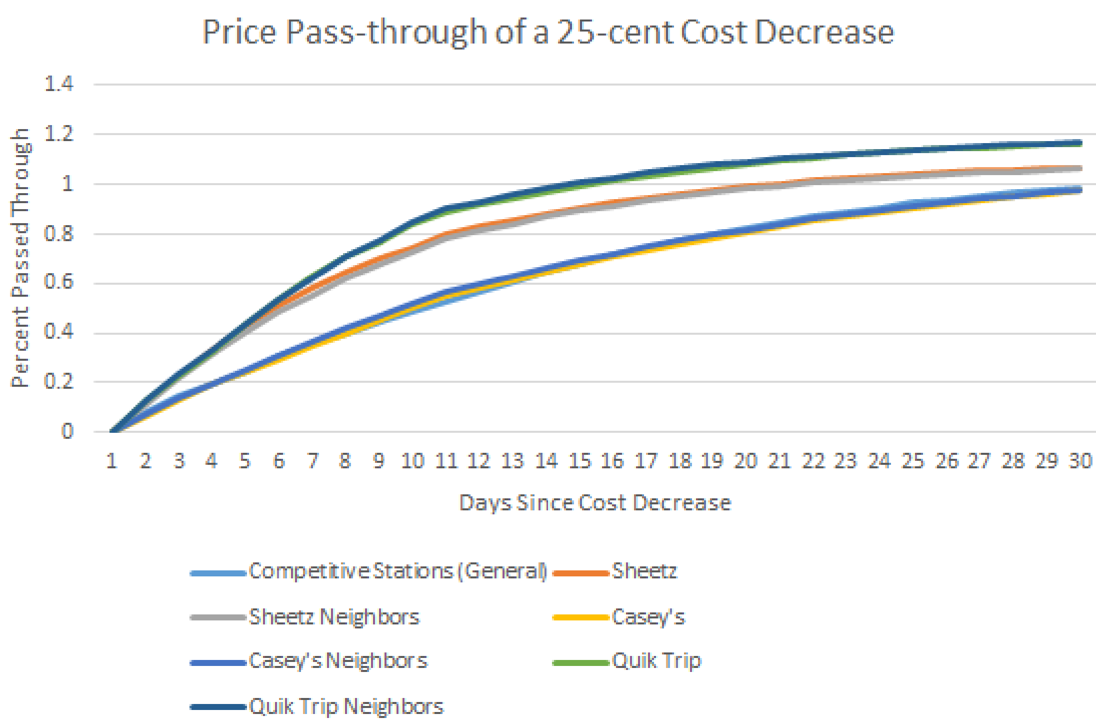

2.5 Conclusion

This paper demonstrates that small, independent or franchised firms in gasoline markets may use well-informed competitors' prices as signals about future costs that are otherwise costly to acquire. My analysis and simulations suggest that in markets that have a QuikTrip or Sheetz location, prices generally rise and fall more quickly following a cost shock than they do in similarly sized markets that do not have one of those three company-operated stations. Markets with a Casey's General Store also respond quickly to cost increases, but they lower prices at about the same rate as typical competitive markets.

These results have implications beyond simple competitive effects, which have been studied extensively in gasoline markets. Previous models do not explain why in the case of a cost increase, a subset of corporate entities' locations consistently increase price earlier than their competitors, and those local competitors increase price more quickly than typical competitive firms. One way to explain this is to use a price coordination model that allows for imperfectly informed firms. Firms that are able to acquire relevant cost information more easily than others tend to change price quickly. Their competitors then observe that price change and adjust their prices according to that station's price instead of acquiring the relevant information directly.

Chapter 3

Lowering the Voting Age for More Votes: Strategic Expansion of the Franchise

3.1 Introduction

In most models of majority rule voting, the pool of voters is taken as exogenous. However, in theory, political entities may be expected to choose to change the voting pool by enfranchising new groups of people. One such group is the pool of younger citizens, who otherwise would have to wait until they reach the minimum voting age to participate in the electoral process. Countries have frequently amended their voting age policy, and in most of those cases, they reduce the minimum age to vote. Many countries reduced the voting age to 18 in the 20th century, and Brazil and some states in Austria have further lowered it to 16. In the United States, where the national voting age is 18, many states have lowered the primary election voting age to 17 for voters who would be 18 for a presidential election. Most other worldwide referenda for changing the voting age have proposed a minimum age lower than that which preceded the proposal.

Two questions naturally arise from this trend. First, why would a country change the voting age? If the current voting pool votes in favor of a particular policy or policy set, then why would any set of decisive voters choose to expand the pool and potentially upset their decisiveness? Second, why is it that voting age shifts are typically realized as a reduction in the minimum age? Why is the minimum never raised, and why is there usually no maximum?

The motivations for these questions revolve around the strategic play by the politicians in power. Most of the literature on the enfranchisement of new voters focuses on the effect of bringing in new voters. Husted and Kenny (1997) showed evidence of a strong relationship between expansion of the franchise to blacks and welfare spending. The rationale used to describe this effect was that black citizens at the time existed at an extreme end of the income distribution; they were relatively much poorer, so when they were added to the voting pool, they favored more redistribution. This conclusion suggests that younger voters, who might also be on the lower end of the income distribution, might similarly favor redistributive policy. Moreover, as long as younger voters have a tendency to fall on one side of a particular issue, their presence in the voting pool could determine the result of an election. Additionally, Lovell (1975) and later Kenny (1978) discussed the general theory behind the shift of the voting franchise to the poorer segment of the population.

Acemoglu and Robinson (2000) created a framework for a strategic expansion of the franchise, but this theory was centered on the possibility of revolution; those in power granted the franchise to the poor in order to avoid being overthrown by those very same previously disenfranchised poor. My approach will be to consider a different threat: that of the current voting pool changing its mind (via realization of preference shocks, etc.), therefore forcing the incumbent to add more voters to the pool who might favor his policies even more than the established median voter.

Llavador and Oxoby (2005) explored enfranchisement of non-elites with respect to economic growth and development. Their work centered on groups of elites served by political parties. Working class wages and landowner capital played a large part in this model, which allowed for de facto political influence as a direct result of wealth. In that model, forward-thinking political parties strategically expanded the franchise to allow for supportive growing social classes to gain de jure power.

In Part 2 of this paper, I attempt to formalize a model that considers a correlation between a voter's age and his policy preferences, which implies an expected gain for some policymakers from adding a group of similarly-aged voters into the voting pool. Additionally, the model attempts to account for the rarity with which the franchise is constricted. Political entities almost never reduce the pool of voters using democratic means in times of peace. I draw from Barzel and Silberberg (1973) as well as Brunk (1980) to establish a rational voter model that places inherent utility on the act of voting itself. If voters place sufficient value on this act, we would expect them to vote against candidates whom they believe likely to remove them from the pool.

In Part 3, I include historical context for the expansion of the voting franchise to younger voters in various parts of the world. While data are limited to large discrete decisions by governing bodies, anecdotal evidence supports the reasoning and results from Section 2. In Part 4, I conclude the study.

3.2 Theory

To investigate this phenomenon, I consider a population of N potential voters, of which $A \subset N$ of these voters are above a current minimum age a_0, and the remaining $N - A$ voters are between the ages of a_1 and a_0, where a_1 represents a potential new minimum voting age. Citizens receive utility as a function of the distance between their favorite policy x_i and the policy implemented by winning candidate j, x_j^p, as well as of indicator D_i, which is equal to 1 if the citizen is in the voting pool and 0 otherwise:

$$U_i j(x_i, D_{it}) = -(x_j^p - x_i)^2 + \alpha D_{it} + v_{jt} \tag{3.1}$$

Where $D_i = 1$ if $a_i \geq a_0$ and is 0 otherwise, and v_j represents a random, i.i.d. preference shock that citizens have in common for candidate j.[1]

The parameter α represents the utility gained from participating in the electoral process. The rationale behind the inclusion of this parameter follows from Barzel and Silberberg (1973) and later Brunk (1980), along with a small literature justifying the act of voting. In the presence of extremely low probabilities that a particular voter is decisive (e.g. a large voting pool), this literature finds that there must be some utility granted to those for the simple act of voting. These voters gain value from performing a "civic duty,"and therefore continue to vote even with a low likelihood of direct influence. This parameter is not new to economic thought, but it is typically omitted from voter utility functions for sound reasons: if everyone receives the same utility from voting, and everyone votes in each period, then this value shouldn't factor into voters' preference rankings. However, under strategic changes in the voting franchise, this parameter may be pertinent.

[1] Relaxing the assumption of common valence terms would complicate the model, but could also make it more realistic. When considering a candidate as I will in this paper who unexpectedly wins an election, it might make sense to allow for heterogeneity in personal preference for a particular candidate.

Voters examine candidates with the knowledge that one candidate may have an incentive (or even a commitment) to constrict the franchise and reduce the voter's utility by α.

Citizen i's policy preference x_i is partly determined by the voter's age, so that

$$x_i = \beta a_i + \mu_i \tag{3.2}$$

where μ_i is an independent and identically distributed random component. In this model, candidates $j \in (1, 2, , J)$ are policy-motivated, must be able to commit to policy proposals[2], are exogenously chosen from the population of citizens, and have a probability of winning which is a function of both their proposed policy set and the random components v_j. In a two-candidate election ($j \in 1, 2$), we can model the probability that Candidate 1 wins as

$$P_1 = Pr(U_{m1} > U_{m2}) \tag{3.3}$$

where U_{mj} represents the utility for the median voter m when candidate j wins. I model this as a two period game in which all payoffs in the second period are discounted by δ, $0 < \delta < 1$. The intuition behind the model is that in many cases, β is nonzero, meaning that age has some kind of (in this case, linear) correlation with policy preferences. β is assumed to be common knowledge to all players within a political structure.

At the start of the first period ($t = 0$), citizens realize their values of a_i and μ_i (and as a result, their values of x_i and D_i). The candidates observe the preferences of the A citizens who are eligible to vote, as well as those of the $N - A$ voters who cannot vote. They observe their own preferences and the distribution of $v = v_1 - v_2$, and then announce policy proposals x_j^p.

It is important to note that the voting age, while a choice variable for the candidate, is not considered to be directly related to x_j^p. It is relevant only in that it affects what policies might be proposed (and implemented) in the future by candidates who respond to the new pool of voters represented by a different distribution of x_i. In this model, the candidate's choice of the voting age is not part of the proposal. This is a separate choice made outside of the policy proposal, and perhaps differs most importantly from x_j^p in that there is no commitment to any voting age. However, as we will see, voters can predict how candidates will choose this variable.

[2]Commitment to policy proposals is important for the model's logical integrity, but it need not refer to true commitment. Much of the logic for strategic responses can be applied ex ante, i.e. a candidate who proposes X and instead attempts to enact Y is unlikely to be re-elected for this reason alone.

After proposals are announced, voting citizens realize the actual values v_j, and proceed to vote for whichever candidate maximizes their long-term expected utility. The median voter, m, determines the winner of this election.

After the result of the election is determined, the winning candidate enacts policy. In this model, a candidate must commit to policy x_j^p, and hence enacts the set of policies on which he campaigned. If the winner chooses not to change the voting age, then the following period ($t = 1$) consists of a second election using the same pool of voters. I will briefly consider the infinite horizon extension of this model for the purposes of potential econometric testing, which will be left to further study with more complete data.

If the winner chose to lower the age to a_1, then the second period confronts players with another election, but with an additional $N - A$ players who are now eligible to vote. These new players clearly influence the median x preference, since they have low values of a_i, which influences their own preferred policies of x_i.

For the two-period model, each candidate will choose an optimal policy proposal x_j^*, somewhere between his favorite policy and the median voter's favored policy. In a setting with parties 1 and 2, with 1 as the party whose policy value is below the median, we have $x_2 > x_2^* > x_m > x_1^* > x_1$. The valence shocks are realized, and voters look forward to the next election. For the purposes of simplifying the model, I assume that the candidate for the losing party in the second election is ideologically identical to the losing candidate in the first election. That is, for an election in which party 1 wins, the loser favored policy x_2. In the second period, the now incumbent party 1 candidate will face a challenger who also favors x_2. This allows for voters to more accurately look forward to the next election.

A key assumption is that the random v value is re-drawn in the second period after the new candidates announce their policy proposals, a reflection of the unobserved difference between the loser in the first election and his same-party replacement in the second election. This implies while the expected median voter's preferences have shifted, a particular voter doesn't know what kind of policy she will be voting for (since she doesn't know how she will respond to the new set of candidates). For this reason, a voter could oppose a change in the voting age, but still vote for a candidate who supports it and against a candidate who opposes it. This is an important implication, because it means that even under the median voter theory, when it seems logical that the median voter would oppose adding in new voters who would surely shift the preference distribution, the voting franchise could expand.

Without the upcoming election, policy proposals can be derived from maximization of win probabilities. For Candidate 2, we have:

$$x_2^* = x_2 - 2(x_2^* - x_m)\frac{f(z)}{F(z)}[(x_1^* - x_2)^2 - (x_2^* - x_2)^2] \tag{3.4}$$

where $z = (x_1^* - x_m)^2 - (x_2^* - x_m)^2$. However, when voters account for the second election and the possibility of a shift in the median, they can see that Candidate 2's favored policy in period 2 will be identical to the equation above, but instead with a new median voter m' in place of m.

If we consider the case in which $\beta > 0$ (so that $m' < m$), the effect on the fraction $\frac{f(z)}{F(z)}$ is ambiguous [3], but the effect on the prior term $-2(x_2 - x_2^*)$ is unambiguously negative. There is a likely downward shift for Candidate 2's proposed policy, away from his own favored policy. Candidate 2 would oppose expanding the franchise in this case. However, in the case in which $\beta < 0$ (so that $m' > m$), then x_2^* in the second election would be closer to x_2, and Candidate 2 would be made better off with the extra voters. He has an incentive to expand the franchise to younger voters who will shift the median and allow him to enact a policy that is closer to his ideal.

Voters take this into account when making their decision in the first election, choosing to maximize discounted lifetime utility. The median voter will be hurt by the change in the voting age; any shift in second-election proposed policy would move away from the median. However, in the case that the first-election winner would have been considered a predicted loser (high first-period realization of v_2), then a shift in the voting franchise might be less of a disadvantage to the median voter than Candidate 2's loss in the first election. To obtain this, we need

$$\delta p U_m(x_2^{'}) + (1-p)U_m(x_1^{'}) < U_m(x_1^*) - U_m(x_2^*) + v \tag{3.5}$$

In the above equation, the left-hand side refers to the expected period 2 loss that the current median voter experiences due to a shift in the voter distribution which results from an expansion of the franchise. Here, p refers to the period 2 probability that Candidate 2 wins re-election, and $x_j^{'}$ refers to the policy proposal by candidate j given a shift in the median voter from m to m'. The right-hand side of the above equation refers to the utility that the median voter obtains from voting

[3] As long as the effect on this term is not strongly positive (more positive than the effect on $-2(x_2 - x_2^*)$ is negative), the results here are unaffected by this fraction.

for Candidate 1 in period 1. This side includes the aggregate valence shock $v = v_2 - v_1$.

We can solve for the necessary value of v, denoted here as $\tilde{v}$, that would cause voters to favor Candidate 2 even when he would be expected to lower the voting age:

$$\tilde{v} = \delta[pU_m(x_2^{'}) + (1-p)U_m(x_1^{'})] - U_m(x_1^*) + U_m(x_2^*) \tag{3.6}$$

We can use this to derive an effect of the change in the probability of a Candidate 2 victory in period 2. It is important to note that this probability is not endogenous to $\tilde{v}$, since $\tilde{v}$ refers to period one valence realizations, and p reflects period 2 valence shocks, which are re-drawn from the valence distribution:

$$\frac{\partial \tilde{v}}{\partial p} = \delta U_m(x_2^{'}) - U_m(x_1^*) \tag{3.7}$$

This implies that as p increases, the gap between the current median voter's utilities from a Candidate 2 victory and a Candidate 1 victory in period 2 has a larger impact on the valence shock needed for Candidate 2 to win in the first period.

This model could be extended to an infinite horizon game, with each period representing an office term. The resulting citizen utility can be formalized as:

$$U_{ij}(x_i, D_i) = \delta^t[-(x_j^p - x_i)^2 + \alpha D_i + v_{ij}] \tag{3.8}$$

so that citizens can anticipate potential proposals for a change in the voting age. However, because of the random preference shocks, there still exist distributions $f(v_j)$ and $g(\mu_i)$ that could result in the election of an official who proposes and passes a change in the voting age. Even aside from these disturbances, a sufficiently large δ paired with a large enough difference between the candidates' proposed policies will result in voters sacrificing future policy gains (through a shift in policy toward the young voters) in favor of current policy gains.

In order for this model to predict only expansions - rather than contractions - of the voting pool, a sufficiently high α would be required to deter voters from allowing a candidate to take away their voting rights. The utility losses from a contraction are not measured only in a shift in the voter distribution, but also in the lack of participation among those voters who have a say now but would not have one later. With a large enough α, those voters could near-unanimously vote against a candidate who would favor such a change in the voting age.[4]

One fact that still remains a mystery according to this model is the tendency for the voting age to shift downward once, and then not shift again. Also, most countries leave their voting age at 18, or at the very youngest, 16. This is not explained by my above model. One way to account for this is by possibly allowing for correlation between age and the preference shock v_j. This would lead to more erratic behavior by younger voters, and risk-averse older voters might be more reluctant to allow for these citizens to participate. In addition, it could be the case that these voters are less predictable in general. If a candidate does not see a clear advantage to adding them into the pool, he will not grant the franchise to a younger age group.

Another solution to this problem can be found in relaxing the assumption of a linear relationship between age and policy preference. A non-linear relationship could lead to more extreme predicted shifts in the voter preference distribution, which would deter candidates from choosing to expand the voting franchise.

This model describes how the voting franchise can change as a result of strategic interaction between candidates and voters. However, it should be noted that the procedure to change the franchise can be far more complicated than the simple act of a single politician. Nevertheless, the logic in my model can be expanded to a broader political regime involving multiple political agents who would gain from the prospect of younger voters participating in elections.

3.3 Historical Context

The United States established its current voting age of 18 when Congress passed the 26th amendment in 1971. The age was lowered from 21 after decades of debate, culminating in a Supreme Court ruling that enabled Congress to lower the national voting age. Congress voted overwhelmingly to pass a bill that would lower the voting age to 18. President Richard Nixon initially opposed the measure from a constitutional standpoint, but said he supported the idea of the 18-year-old voter. If we consider younger American voters to lean more liberal due to issues such as income redistribution, social progressivism, and draft status, then this story anecdotally fits the model created in Section 2 of this paper. A highly Democratic Congress favored the introduction of younger voters, which was followed by even more intensely Democratic Congresses over the next two sessions. Meanwhile,

[4]This phenomenon would also account for the lack of a maximum voting age in any observable setting.

the Republican president was less enthusiastic about a change; he followed two terms of Democratic presidency, and was seeking re-election. It would not be in his interest to support a voting age change. Ultimately, the decision was made by Congress, and was not vetoed by Nixon. This could be explained by Nixon's reasoning before the Supreme Court ruling for why he opposed the legislation; he had said he supported the idea, so now that it was feasible, it may have been less politically feasible to go back on that promise than to simply allow for an influx of opposition to the voting pool. He would instead have to moderate his own policies and policy proposals in order to bid for re-election.

Brazil lowered its minimum voting age from 18 to 16 in 1888, following 21 years of military rule. The move from oppression to democracy is indicative of a shift in power from right to left, which may support the idea that the change was strategic. The new regime, in order to maintain power, added voters who were likely to support it in the future.

In 2007, Iranian president Mahmoud Ahmadinejad suggested that the voting age be reduced from 18 to 15. The Ahmadinejad regime had been known for some extremely conservative views, so it may at first seem to run counter to the specified model. However, a 2013 survey from InterMedia[5] found that the Iranian youth is considerably more conservative than other young voters. About 48 percent of the Iranian youth fell into the ultra-conservative or conservative categories, while only 18 percent fell into the liberal "Non-Traditionalist" category. As such, President Ahmadinejad had a clear incentive to bring these youth into the voting pool in order to push his policies further to that extreme. The example of Iran illustrates that a strategic choice model of lowering the voting age explains such a policy regardless of political affiliation. In contrast, any theory that supposes that more progressive politicians are naturally inclined to lower the voting age for the sake of inclusion cannot account for the Iranian change in policy.

While data on voting age changes - specifically with respect to the political climate at the time - are limited, there exists some anecdotal evidence that the voting age may be a strategic move similar to the way in which I have modeled it. The infrequency with which this change takes place indicates that populations tend to resist this change, and that governments are only willing to expand the franchise when it is clearly more advantageous than simply attempting to get re-elected by the same voting pool.

[5]InterMedia is an independent non-profit research center which aims to collect data on political ideology and mobilization.

3.4 Conclusion

While previous literature has focused mainly on the effects of an expansion of the voting franchise, this paper has considered the possibility that such expansion is a strategic decision on the part of the policy maker and should thus be considered an endogenous choice. My specification involved a one-dimensional policy spectrum with probabilistic voting, which leads to some leaders being unexpectedly elected despite a lack of popularity for their policy proposals per se. Such a politician would have an incentive to change the rules of the election for the next period, since he doesn't believe he can win on his own policy merits again. Expanding the franchise to voters closer to himself in ideology could sway the next election favorably. Additionally, even if the incumbent believes he will win a re-election bid, an expansion of the franchise to like-minded citizens would allow him to run on policies that are closer to his own ideal; he can avoid having to run on a moderate platform to appease the old median voter.

The relevant historical information on voting age changes is sparse, but what is available somewhat supports the expansion of the franchise to the young as an endogenous choice. At the very least, it fails to refute the argument presented here. It is reasonable to suspect that regimes have expanded the franchise strategically in order to gain support from those new voters in the future.

Appendices

Appendix A Observed Frequencies at Which Gasoline Stations Change Price

Table 1: Frequencies for price increases

Market Type	Obs	Pr(Inc)	Std. Err.
All Stations	8955	.0715	.0006
Isolated	364	.0579	.0038
Pairs	206	.0576	.0036
Crowded	1313	.0701	.0014

Table 2: Frequencies for price decreases

Market Type	Obs	Pr(Inc)	Std. Err.
All Stations	8955	.1443	.0012
Isolated	364	.0784	.0051
Pairs	206	.0852	.005
Crowded	1313	.1458	.0027

Appendix B Coefficient Estimates For Gasoline Station Price Responses

Table 3: Price Increase Magnitudes

Market Type	Obs	Mean Increase	Std. Err.
All Stations	169859	.0715	.0006
Isolated	3958	.0579	.0038
Pairs	3235	.0576	.0036
Crowded	24157	.0701	.0014

Table 4: Price Decrease Magnitudes

Market Type	Obs	Mean Decrease	Std. Err.
All Stations	343315	-.0445	.0001
Isolated	5951	-.0813	.0012
Pairs	4969	-.0646	.0011
Crowded	50572	-.0442	.0003

Table 5: Results for Equation 5: Price Increase Likelihood

Coefficient	Isolated	Pairs	Crowded
λ^{IL}	-4.393*** (0.289)	-3.154*** (0.333)	-4.569*** (0.133)
λ^{IH}	-3.035*** (0.573)	-2.194*** (0.520)	-3.254*** (0.272)
τ^{I}	0.244*** (0.081)	0.3326*** (0.083)	0.340*** (0.030)
γ_1^I	0.842 (0.583)	1.543* (0.599)	2.149*** (0.217)
γ_2^I	2.964*** (0.579)	2.796*** (0.592)	3.618*** (0.215)
γ_3^I	3.844*** (0.583)	2.848*** (0.604)	4.038*** (0.217)
γ_4^I	2.474*** (0.608)	2.965*** (0.627)	3.155*** (0.227)
γ_5^I	3.524*** (0.606)	3.603*** (0.616)	4.618*** (0.226)
γ_6^I	5.235*** (0.597)	4.550*** (0.612)	5.710*** (0.224)
γ_7^I	3.931*** (0.595)	2.448*** (0.607)	3.505*** (0.224)
γ_8^I	2.786*** (0.608)	1.073*** (0.621)	2.827*** (0.229)
γ_9^I	0.682 (0.621)	-0.680 (0.630)	0.646*** (0.230)
γ_{10}^I	-0.468 (0.629)	-0.908 (0.636)	-0.222 (0.231)
γ_0^I	-3.307*** (0.068)	-3.023*** (0.068)	-2.946*** (0.025)

Notes: λ^{IL} represents a firm's price increase probability response, given $\mu_{jt} < 0$. λ^{IH} represents a firm's price increase probability response, given $\mu_{jt} > 0$. $\gamma_1^I, ...\gamma_{10}^I$ represent the lagged responses to cost shifts. γ_0^I represents the constant term. τ^I is the additional price increase likelihood for the indicator for $\mu_{jt} < 0$.

Table 6: Results for Equation 6: Price Decrease Likelihood

Coefficient	Isolated	Pairs	Crowded
λ^{DL}	3.404*** (0.507)	2.476*** (0.483)	2.306*** (0.172)
λ^{DH}	1.875*** (0.195)	1.392*** (0.204)	1.873*** (0.086)
τ^{D}	-0.236*** (0.062)	-0.316*** (0.064)	-0.295*** (0.019)
γ_1^D	-0.204 (0.468)	-0.869* (0.482)	-0.981*** (0.154)
γ_2^D	-2.525*** (0.479)	-2.56*** (0.494)	-2.867*** (0.158)
γ_3^D	-2.034*** (0.482)	-2.363*** (0.495)	-3.482*** (0.159)
γ_4^D	-1.471*** (0.494)	-1.607*** (0.504)	-2.241*** (0.162)
γ_5^D	-0.798* (0.486)	-0.57 (0.494)	-1.191*** (0.159)
γ_6^D	-2.306*** (0.475)	-2.244*** (0.484)	-2.13*** (0.156)
γ_7^D	-2.855*** (0.467)	-3.222*** (0.478)	-2.972*** (0.154)
γ_8^D	-1.863*** (0.461)	-2.324*** (0.474)	-1.92*** (0.152)
γ_9^D	-2.804*** (0.452)	-2.046*** (0.466)	-1.73*** (0.15)
γ_{10}^D	-3.208*** (0.463)	-2.345*** (0.474)	-2.126*** (0.152)
γ_0^D	-2.408*** (0.035)	-2.067*** (0.037)	-1.611*** (0.011)

Notes: λ^{DL} represents a firm's price decrease probability response, given $\mu_{jt} < 0$. λ^{DH} represents a firm's price decrease probability response, given $\mu_{jt} > 0$. $\gamma_1^D, ...\gamma_{10}^D$ represent the lagged responses to cost shifts. γ_0^D represents the constant term. τ^D is the additional price decrease likelihood for the indicator for $\mu_{jt} < 0$.

Table 7: Results for Equation 7: Price Increase Magnitude

Coefficient	Isolated	Pairs	Crowded
θ^{IL}	-0.133*** (0.023)	-0.140*** (0.02)	-0.172***(.008)
θ^{IH}	-0.026 (0.037)	0.091*** (0.029)	-0.01 (0.015)
ν^{I}	-0.008 (0.005)	-0.004 (0.005)	-0.005*** (0.002)
β^I_1	0 (0.033)	0.01 (0.029)	-0.002 (0.011)
β^I_2	0.049 (0.037)	0.069** (0.031)	0.017 (0.011)
β^I_3	0.049 (0.042)	0.055 (0.034)	0.043*** (0.013)
β^I_4	-0.024 (0.043)	0.047 (0.036)	0.048*** (0.014)
β^I_5	-0.047 (0.04)	-0.025 (0.034)	0.048*** (0.013)
β^I_6	0.064* (0.037)	-0.033 (0.031)	0.033*** (0.012)
β^I_7	0.014 (0.036)	0.04 (0.032)	0.022 (0.012)
β^I_8	0.049 (0.037)	0.038 (0.031)	0.007 (0.012)
β^I_9	-0.013 (0.041)	-0.005 (0.034)	-0.019 (0.013)
β^I_{10}	-0.152*** (0.045)	0.012 (0.037)	-0.036** (0.014)
β^I_0	0.087*** (0.004)	0.063*** (0.004)	0.06*** (0.001)

Notes: θ^{IL} represents a firm's price increase magnitude response, given $\mu_{jt} < 0$. θ^{IH} represents a firm's price increase probability response, given $\mu_{jt} > 0$. $\beta^I_1, ...\beta^I_{10}$ represent the lagged responses to cost shifts. β^I_0 represents the constant term. ν^I is the additional price increase likelihood for the indicator for $\mu_{jt} < 0$.

Table 8: Results for Equation 8: Price Decrease Magnitude

Coefficient	Isolated	Pairs	Crowded
θ^{DL}	-0.054* (0.028)	-0.039** (0.02)	-0.037*** (0.006)
θ^{DH}	-0.221*** (0.011)	-0.185*** (0.009)	-0.211*** (0.003)
ν^{D}	-0.007** (0.003)	-0.009*** (0.003)	-0.009*** (0.001)
β_1^D	-0.09*** (0.022)	-0.039** (0.018)	-0.055*** (0.005)
β_2^D	-0.088*** (0.024)	-0.053** (0.02)	-0.069*** (0.005)
β_3^D	-0.035 (0.025)	-0.068*** (0.021)	-0.06*** (0.006)
β_4^D	-0.048* (0.028)	-0.014 (0.022)	-0.042*** (0.006)
β_5^D	-0.061** (0.026)	-0.019 (0.021)	-0.034*** (0.006)
β_6^D	-0.104*** (0.024)	-0.02 (0.02)	-0.053*** (0.005)
β_7^D	-0.065*** (0.023)	-0.025 (0.018)	-0.067*** (0.005)
β_8^D	-0.032 (0.021)	-0.042** (0.018)	-0.046*** (0.005)
β_9^D	-0.009 (0.022)	-0.008 (0.018)	-0.037*** (0.005)
β_{10}^D	-0.024 (0.023)	0.03 (0.019)	-0.017*** (0.005)
β_0^D	-0.054*** (0.002)	-0.041*** (0.002)	-0.027*** (0.0004)

Notes: θ^{DL} represents a firm's price decrease magnitude response, given $\mu_{jt} < 0$. θ^{DH} represents a firm's price decrease probability response, given $\mu_{jt} > 0$. $\beta_1^D, ...\beta_{10}^D$ represent the lagged responses to cost shifts. β_0^D represents the constant term. ν^D is the additional price decrease likelihood for the indicator for $\mu_{jt} < 0$.

Bibliography

[1] Acemoglu, D., Robinson, J.A. (2000). Why did the West Expand the Franchise? Democracy, Inequality, and Growth in Historical Perspective. The Quarterly Journal of Economics, Nov 2000, 1167-1199.

[2] Bacon, R.W. (1991) Rockets and Feathers: The Asymmetric Speed of Adjustment of UK Retail Gasoline Prices to Cost Changes. Energy Economics, Vol. 13, pp. 211218.

[3] Barzel, Y., Silberberg, E. (1973) Is the Act of Voting Rational? Public Choice, Fall 1973, 51-58

[4] Borenstein, S., Cameron, A.C., Gilbert, R. (1997). Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes? The Quarterly Journal of Economics 112 (1): 305-339.

[5] Borenstein, S., Shepard, A., (2002). Sticky prices, inventories, and market power in wholesale gasoline markets. RAND Journal of Economics, 33, 2002, 116139

[6] Bresnahan, T., Reiss, P. (1991). The Journal of Political Economy, Vol. 99, No. 5., pp. 977-1009.

[7] Brunk, G.G. (1980) The Impact of Rational Participation Models on Voting Attitudes. Public Choice, 1980, 549-564

[8] Colby, D.C. (1986) The Voting Rights Act and Black Registration in Mississippi. Publius, Autumn 1986, 123-137

[9] Davis, M., Hamilton, J., (2003). Why are prices sticky? The dynamics of wholesale gasoline markets. Journal of Money, Credit, and Banking, 36, 17-37.

[10] Dixit, A., 1991. Analytical Approximations in Models of Hysteresis. The Review of Economic Studies 58, 141-151.

[11] Hastings, J.S. (2004) Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California. American Economic Review, 94(1), 317-328.

[12] Husted, T.A., Kenny, L.W. (1997) The Effect of the Expansion of the Voting Franchise on the Size of Government. Journal of Political Economy, Feb. 1997, 54-82.

[13] InterMedia. (2013) Youth in Iran: A Story Half Told. Young Publics research Paper Series, May 2013.

[14] Karrenbrock, J.D. (1991) The Behavior of Retail Gasoline Prices: Symmetric or Not Federal Reserve Bank of St. Louis Review, July/August, pp. 1929.

[15] Kenny, L.W. (1978). The Collective Allocation of Commodities in a Democratic Society: A Generalization. Public Choice, 1978, 117-120.

[16] Llavador, H., Oxoby, R.J. Partisan Competition, Growth, and the Franchise. The Quarterly Journal of Economics, Aug. 2005, 1155-1189.

[17] Lewis, M. (2011). Assymetric price adjustment and consumer search: An examination of the retail gasoline market.

[18] Lewis, M. (2012). International Journal of Industrial Organization, 30 (4), 342-351.

[19] Nayyar, P.R. (1990) Information Asymmetries: A Source of Competitive Advantage for Diversified Service Firms. Strategic Management Journal, 11, 513-519.

[20] Scherer, F. (1967). Focal Point Pricing and Conscious Parallelism. Antitrust Bulletin 12(2), 495-504.

[21] Shepard, A. (1990) Pricing Behavior and Vertical Contracts in Retail Markets. The American Economic Review, 80(2), 427-431.

[22] Shepard, A. (1993) Contractual Form, Retail Price, and Asset Characteristics. RAND Journal of Economics, 24(1), 5877.

[23] Tappata, M., (2009). Rockets and Feathers: Understanding Asymmetric Pricing, RAND Journal of Economics, 40, 673687.

[24] Vincenty, T. (1975) Direct and Inverse Solutions of Geodesics on the Ellipsoid with Application of Nested Equations. Survey Review, 32 (176), 88-93.

[25] Yang, H., Ye, L., (2008). Search with learning: understanding asymmetric price adjustments. RAND Journal of Economics, 39, 547-564