Price Fluctuations and Market Power: Evidence from Retail Gasoline Markets

Peter D. Lorenz

September 29, 2016

Abstract

This paper examines cost pass-through in the retail gasoline market, focusing on stations with high market power. A large literature has identified pricing patterns in retail gasoline markets, ranging from price asymmetry to focal point tacit collusion. While a small portion of this literature has indicated the tendency for wholesalers with high market power to change price less frequently, this paper aims to firmly establish a link between market power and price volatility at the retail level. In addition, most of this work assumes smooth price adjustments by individual stations, an assumption I observe to be false. Using an amended model that accounts for discrete price adjustments, I discover that geographically isolated stations respond more slowly to cost shocks than do stations with many competitors. Further, I find that isolated pairs of stations competing only with one another raise prices even more slowly than isolated single stations.

1 Introduction

There exists a significant literature explaining pricing patterns in the retail gasoline market. Many of these papers focus on the peculiar habits of stations with many competitors; analyses of price adjustment asymmetry (retail gasoline prices rise more rapidly than they fall), tacit collusion, and price dispersion are extensive. However, the literature is largely lacking in analysis of stations with few competitors. This relative gap in the literature leaves open the question of how firms behave in the absence of widespread competition. In addition, the problem of price stickiness has not been adequately addressed at the retail level. Early observation shows that the average gas station changes its price from one day to the next only 21 percent of the time, despite daily publicly posted changes to wholesale prices. The failure for stations to appropriately adjust for cost shocks indicates that there may be some form of "menu cost"causing stations to hesitate on a price change. No matter the cause, known price rigidity necessarily affects my approach in analyzing price changes.

This paper aims to establish a link between market power and price volatility and pass through. While firms with market power would be expected to provide goods at an inefficiently low level, we would typically assume that such firms would quickly pass cost changes onto consumers. If firms with high market power respond especially slowly to cost shocks, then the efficiency gains from increased competition might be presently understated.

At the retail level, there exists a relationship between the level of a station's local competition and the frequency with which that station changes its price. In general, the more competing stations there are nearby, the more often a station changes its price. The most interesting exception is the case of pseudo-duopolistic markets in which two stations are close to one another, but very far from the next closest station. These stations change their prices even less frequently than isolated single stations.

However, it is not the case that given a price change, a station with fewer neighbors tends to enact a smaller price change; there is no clear relationship between the size of a price change and a station's level of competition. Thus, it is not immediately clear that less competitive stations pass cost changes onto consumers more slowly. Furthermore, in order to make any inference from these observed phenomena, costs must be considered. We would expect a larger cost shock to affect an isolated station differently than it would affect a more competitive station. Monopoly markups imply a larger, faster response to a cost shock. Are isolated stations really slower to change price, or do they simply respond immediately to large cost shocks and adjust fully, leaving no need to change price more frequently in the following weeks? I address this by comparing the price of gasoline at the time of import (a proxy for cost) to the price at which a station sells its gasoline in order to approximate a station's retail markup. I then use this estimated markup to measure the impact of a cost change on the price change decision.

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.

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. Given what is known about the difference in firm behavior between rising and falling prices, my analysis must distinguish between price increases and price decreases.

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 explanation put forth by Lewis and Marvel (2011), as well as a tacit collusion explanation put forth by Borenstein, Cameron, and Gilbert (1997). 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 fail to specify the exact number of firms in competition with one another, which is something I can accomplish with my data.

Borenstein and Shepard (2002) detail a model in which gasoline wholesalers respond gradually to the price of crude oil. In this model, wholesalers incur production level adjustment costs that limit their ability to fully adjust to the price of crude oil. As a result, a single large change in the price of crude oil will be spread out over several periods. They compare this model to a menu cost model, which "focuses on the differences between the market-*clearing* price and the price at which spot transactions actually occur."That paper also predicts that under supply adjustment constraints, wholesalers in areas with more market power will respond differently to cost shocks than those in areas with lower market power. The model, however, does not predict the direction of this effect. That is, it does not specify whether market power will induce more or less stickiness in price adjustment. The authors find empirically that wholesalers in high-market-power areas are slower to respond to prices. My study differs from Borenstein and Shepard's paper in that I focus on retailers directly, and differentiate more precisely between stations with high market power and stations with less market power.

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 the Dixit 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 higher-than-normal stickiness in monopoly 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.

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 share some 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, given a level of demand.

Stations with many nearby competitors, henceforth referred to as "crowded" stations, are likely to have high cross-price elasticities with multiple competitors. 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. The consumer search explanation for this phenomenon has implications on how market power affects price pass-through. The reasoning for the theory is as follows: 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 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 taken hold as an 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 (Scherer, 1967), or a price that ends in the digit "9" (Lewis,

2015) In this way, these stations can keep prices higher than the competitive price, splitting total profits close to what a monopoly would earn. When costs rise, the competing stations raise price quickly, finding a high collusive price that keeps margins high. 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. It is not clear, however, how well these stations are able to coordinate when prices are rising. 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. This coordination problem does not have a clear theoretical result for price rigidity. 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 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. Any failure to do so would indicate some kind of "menu cost," an administrative cost associated with the act itself of changing price. 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 face the above dilemma if immediate coordination is difficult. In measuring general price rigidity, I will also be able to test how isolated station pairs respond to this dilemma. If they are able to coordinate easily, I would expect that they would raise price quickly. If coordination is difficult, I would expect that they would remain at pre-shock prices in the short term after a cost increase.

For crowded stations, a tacit collusion model is again 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.

4 Data and Empirical Approach

I use 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 use the New York Harbor conventional gasoline regular spot price as a proxy for the underlying cost of gasoline for stations in Pennsylvania. I use the U.S. Gulf Coast spot price for stations in the other four selected states. While local wholesale terminals act as intermediaries between harbor and retailer, these wholesalers do not always post a daily price, and when they do, that price is an estimate based on the harbor price. Thus, these harbor prices will adequately serve as a cost signal for individual stations.

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 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.

Table 1: Probabilities 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: Probabilities 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

Tables 3 and 4 describe the average magnitude of price increases and decreases for each market, given a price change $\neq 0$.

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

Notes:

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.

I focus my analysis on stations' response to a particular profit margin, defined as $Margin_{jt} \equiv p_{j,t-1} - c_{jt}$, where $p_{j,t-1}$ is station j's price on day t-1, and c_{jt} is the local wholesale price for station j, plus the gasoline tax for station j's home state as well as the federal gasoline tax.

This variable on its own is vulnerable to several issues in the data. First, my measure of cost is not a precise measurement of the actual cost to an individual station. It does not account for transport cost, nor does it account for wholesaler markups. Second, less competitive stations are expected to have higher markups due to basic spacial differentiation pricing strategy. In order to adjust for these problems, I use a two-step regression approach to link price responses to cost shocks, while controlling for idiosyncratic station behavior and conditions. In the first step, I regress price against cost using the following fixed effects model:

$$p_{jt} = \alpha_j + \phi c_{jt} + \mu_{jt} \tag{1}$$

From Engle and Granger (1987), I can use an error correction model in which I use the residual μ_{jt} from (1) as an estimate for a station's deviation from its expected markup. Using my result from (7), I model price changes in the following way:

$$\Delta p_{jt} = \sum_{s=1}^{10} \beta_s \Delta c_{t-s} + \theta(\hat{p}_{jt} - \alpha_j - \phi c_{jt}) \tag{2}$$

or:

$$\Delta p_{jt} = \sum_{s=1}^{10} \beta_s \Delta c_{t-s} + \theta \mu_{jt} \tag{3}$$

where $\beta_s \Delta c_{t-s}$ represents a lagged cost response term for s days prior to day t. These effects statistically drop to zero after about 10 days in all cases, so I have restricted the lagged cost responses to 10 days.

While μ_{jt} and Δp_{jt} are clearly correlated, I experience heavy interference with small price adjustments along the entire range of μjt. For example, after a large cost increase, a station may engage in a series of price increases. If one of them was larger than necessary, the station might decrease its price for a few days before continuing in the series of price increases. To avoid this problem, I follow 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:

$$I^H = \begin{cases} 1 & \text{if } \mu_{jt} > 0 \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

Following this approach, I have

$$\Delta p^I_{jt} = \beta^L_0 + \sum_{s=1}^{10} \beta^L_s \Delta c_{t-s} + \theta^L \mu_{jt} + \varepsilon_{jt} \tag{5}$$

for price increases, and

$$\Delta p^D_{jt} = \beta^H_0 + \sum_{s=1}^{10} \beta^H_s \Delta c_{t-s} + \psi I^H + \theta^H \mu_{jt} + \varepsilon_{jt} \tag{6}$$

for price decreases, where θ^L refers to the firm's price response when $I^H = 0$, and θ^H refers to its price response when $I^H = 1$. Using similar reasoning, I develop a logistic probability model for the binary decision to change price at all.

We have:

$$Pr(Inc) = \gamma^L + \sum_{s=1}^{10} \gamma^L_s \Delta c_{t-s} + \lambda^L \mu_{jt} + \epsilon_{jt} \tag{7}$$

for price increases, and

$$Pr(Dec) = \gamma^H + \sum_{s=1}^{10} \gamma^H_s \Delta c_{t-s} + \tau I^H + \lambda^H \mu_{jt} + \epsilon_{jt} \tag{8}$$

for price decreases.

5 Results

Using the 2-step estimation technique described above, I record the relevant parameter estimates in Table 5, separated by market type:

Table 5: Results for Equations 5-8

Coefficient	Isolated	Pairs	Crowded
θ^L	-0.133*** (0.023)	-0.140*** (0.02)	-0.172***(.008)
θ^H	-0.221*** (0.011)	-0.185*** (0.009)	-0.211***(0.003)
λ^L	-4.39*** (0.289)	-3.15*** (0.333)	-4.57***(.133)
λ^H	1.87***(.194)	1.39***(0.204)	1.87***(0.0141)

Notes: θ represents a firm's price magnitude response. λ represents a firm's price change probability response. β_0 refers to the constant term on the price magnitude regression. γ_0 refers to the constant term on the price change probability regression.

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. 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 significantly 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

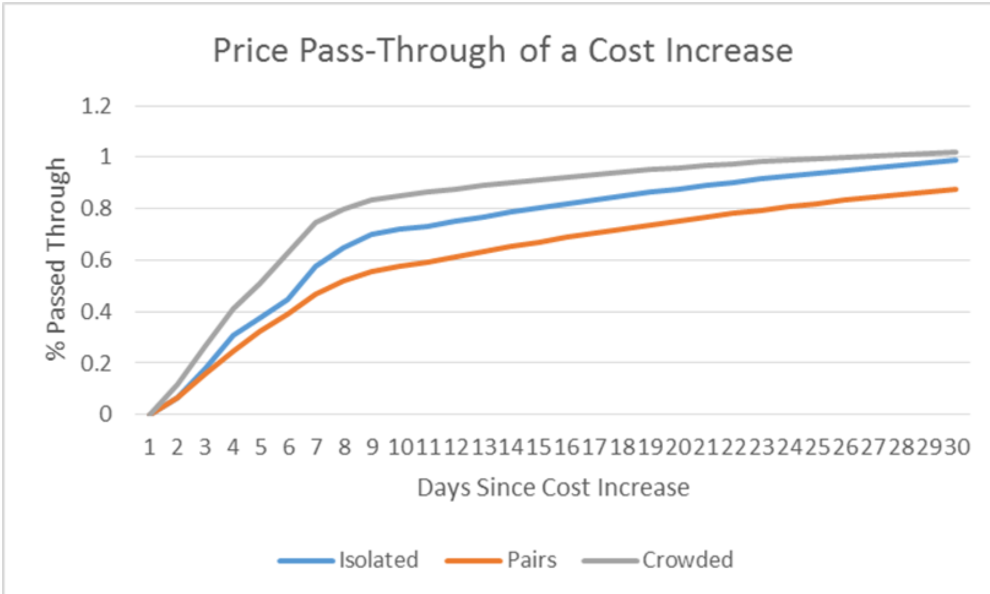

Figure 2 shows estimated price paths after a price decrease. Crowded stations again exhibit the speediest price adjustment, reaching 80 percent pass-through by day 18. Note that this is considerably slower than the upward pass-through for the same market type. This is consistent with previous work on 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 2

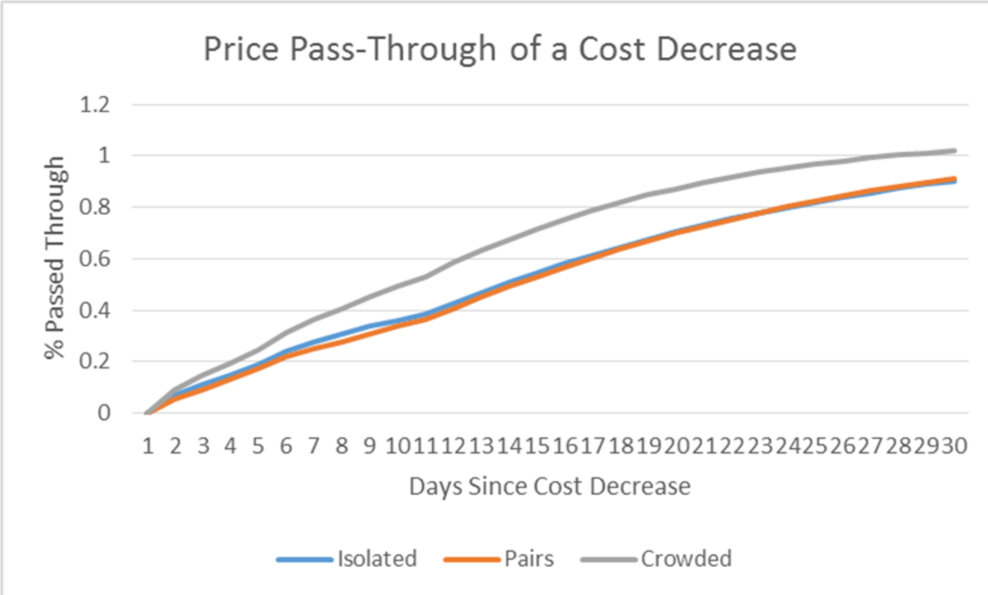

This comparison between isolated stations and station pairs supports the claim that retail gasoline stations use focal point prices to engage in tacit collusion. Specifically, it supports the claim that they use the most recently posted price as a point of collusion. When prices are falling, these stations 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.

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, opting to avoid the penalty for being more expensive than the nearest competitor.

7 Appendix A

Table 6: Additional Results for Equations 5-8

Coefficient	Isolated	Pairs	Crowded
β_0^L	0.0865*** (0.004)	0.0625*** (0.004)	0.0601*** (0.001)
β_0^H	-0.054*** (0.002)	-0.0409*** (0.0016)	-0.0269*** (0.0004)
γ_0^L	-3.31*** (0.068)	-3.02*** (0.068)	-2.95*** (0.025)
γ_0^H	-2.41*** (0.004)	-2.07***(0.037)	-1.61*** (0.011)

Notes: β_0 refers to the constant term on the price magnitude regression. γ_0 refers to the constant term on the price change probability regression.