UC San Diego

Recent Work

Title

The Speed of Gasoline Price Response in Markets With and Without Edgeworth Cycles

Permalink

<https://escholarship.org/uc/item/8j36j1s1>

Authors

Lewis, Matt
Noel, Michael

Publication Date

2009-03-03

THE SPEED OF GASOLINE PRICE RESPONSE IN MARKETS WITH AND WITHOUT EDGEWORTH CYCLES

Matthew Lewis
The Ohio State University*

Michael Noel
University of California
San Diego[†]

March 3, 2009

Abstract

Retail gasoline prices are known to respond fairly slowly to wholesale price changes. This does not appear to be true for markets with Edgeworth price cycles. Recently, many retail gasoline markets in the midwestern U.S. and in other countries have been shown to exhibit price cycles, in which competition generates rapid cyclical retail price movements. We show that cost changes in cycling markets are passed on 2 to 3 times faster than in markets without cycles. We argue that the constant price movement inherent within the Edgeworth cycle eliminates price frictions and allows firms to pass on cost fluctuations more easily.

*Department of Economics, E-mail: mlewis@econ.ohio-state.edu.

[†]Department of Economics, E-mail: mdnoel@ucsd.edu

Introduction

Economists study the speed with which cost changes are passed through to prices in order to better understand how efficiently markets are working. Sticky prices and slow passthrough have important implications for macroeconomic policy. On a more micro level, studying passthrough can reveal how imperfect competition is affecting market outcomes.

In gasoline markets, the empirical literature has shown that retail prices tend to respond slowly and asymmetrically to wholesale costs, with cost increases passed along more quickly than cost decreases. Several explanations have been raised for these patterns. Lewis (2005) suggests that imperfect information and consumer search contribute to slow and asymmetric price response. Verlinda (2008) and Deltas (2008) provide evidence that greater local market power is associated with more asymmetric price response. Overall there is little consensus as to which market fundamentals most strongly influence the extent or asymmetry of delayed gasoline price passthrough.

Several recent studies have also identified that some retail gasoline markets exhibit frequent retail price cycles. These cycles occur independent of cost changes and strongly resemble the Edgeworth price cycle equilibria formalized by Maskin and Tirole (1988). Average retail prices in these cycling markets tend to periodically jump (often 10 cents/gallon or more) within one or two days, and then fall gradually. As soon as falling prices become low relative to wholesale costs, prices jump again and another cycle begins, with a typical cycle lasting one to two weeks. These cycles have been empirically documented in cities within the midwestern U.S. [Lewis (Forthcoming); Doyle et al. (2008)], Canada [Eckert (2003); Noel (2007a); Noel (2007b)], and Australia [Wang (2008)]. There is some evidence that the existence of these cycles may be related to the market structure and concentration of retailers [see Noel (2007a); Lewis (Forthcoming); Doyle et al. (2008)], but the impact that cycles have on the competitiveness of the local market is still not well understood.

In this paper we study how Edgeworth cycles impact the speed with which retail gasoline prices respond to wholesale cost changes. Day to day retail price movements are larger and much more frequent in cities that exhibit Edgeworth cycles than those that do not. More importantly, retail price movements are driven not only by wholesale cost changes (as in a typical market)

but also by an independent mechanism of cyclical price fluctuation generated by retail competition. We hypothesize that this environment of continually changing prices eliminates frictions and allows cost changes to be passed through more quickly to the retail price.

In most cases, existing studies use distributed lag models to describe retail price response patterns. These models work well for measuring the speed of response in non-cycling markets (particularly at the market level). However, in markets that exhibit price cycles, distributed lag models are unable to capture the large and periodic changes in retail margins. Therefore, we estimate price response in cycling markets using a Markov switching regression framework that incorporates the Edgeworth cycle price dynamics. By accurately modeling cyclical pricing behavior we are able to identify how a cost change affects the shape of the current price cycle and why this influences the speed of cost passthrough.

Using panel data from U.S. cities with and without Edgeworth price cycles, we show that prices in markets without cycles respond much more slowly to wholesale cost fluctuations than in cities with cycles. On average, prices in non-cycling markets take three to six weeks to fully reflect a change in costs, while cost changes in cycling markets are fully incorporated into the retail price after only five or six days. Moreover, cities with cycles have much faster passthrough even after controlling for differences in market structure that could influence the speed of price response. In other words, quicker price response in cycling markets appears to be generated by the price cycles themselves rather than by the underlying market structure characteristics that may make cycles more likely to occur. This result is one of the first to explicitly identify how the nature of local retail competition can influence the speed of price response. In addition, within the Edgeworth cycle literature this study provides one of the most concrete measures of how cycles affect market performance and competitiveness.

Edgeworth Price Cycles in Retail Gasoline Markets

The Maskin and Tirole (1988) model contains two identical homogeneous product firms competing in an alternating move game. They show that one equilibrium of the game results when each firm responds by slightly undercutting the other firm's price from the previous period in order to steal the market demand until the other firm can lower its price again. Price continues to fall until

profit margins are zero, at which point one firm raises its price and the cycle of undercutting begins again.¹ Although this model is highly stylized, the predicted pattern of prices strongly resembles the gasoline price movements observed in some markets. Furthermore, the alternating move nature of the game resembles the back-and-forth strategic price jockeying that is often associated with competition between neighboring gas stations. Noel (2008) also shows that Edgeworth cycle equilibria do exist in a number of extensions to the basic Maskin and Tirole setting, including cases with differentiated firms and with more than two firms.

Unlike previous studies of Edgeworth cycles in gasoline markets, we use a broad panel of daily retail and wholesale gasoline prices from 90 cities during 2004 and 2005. The sample represents most of the major cities in the midwestern US and from nearby states in the South and Mid-Atlantic regions. Average retail prices for each city are provided by the Oil Price Information Service (OPIS) based on a survey of gasoline stations in each market.² The wholesale prices of *unbranded* gasoline at the local distribution terminal (called the *rack*) are used as the measure of the local wholesale cost of gasoline.³ The cross section contains many cities that exhibit retail cycles and many that do not. We use this feature to empirically identify systematic differences in the manners with which retail prices in these different types of cities respond to wholesale price fluctuations.

Identifying Markets with Retail Price Cycles

Analyzing differences in wholesale cost passthrough between prices in cities with and without Edgeworth price cycles requires a method of classifying the cities among these two groups. In many cases it is clear even from casual observation whether a city exhibits retail price cycles. Figure 1 shows the typical retail and wholesale price movements from two of the sampled cities: Pittsburgh, PA and Columbus, OH. Though the wholesale prices for the two cities are fairly similar, the average retail price in Columbus moves in a highly cyclical pattern that is largely indepen-

¹Once profit margins are zero ($P=MC$), both firms would like price to increase, but they want the other firm to be the one who raises the price first. As a result, a war of attrition (ie. mixed strategy response) occurs until one of the firms raises price, allowing the other firm to undercut and earn profits.

²Retail prices are for a gallon of regular (87 octane) gasoline with all relevant taxes removed.

³Rack prices are also provided by OPIS. The unbranded gasoline rack price represents the best available measure of the opportunity cost of retail gasoline in a particular city. Rack prices for *branded* gasoline sold by large refining companies to their branded dealers are not a good measure of the true cost of gasoline because refining companies manipulate their branded rack prices in order to extract extra retail profits from their dealers.

Figure 1: Typical Gasoline Prices in Columbus and Pittsburgh.

dent of wholesale price movements. In contrast, Pittsburgh’s retail price shows much less day to day fluctuation and most of the movement in the retail price appears to correspond with changes in wholesale prices.

Lewis (Forthcoming) proposes a more systematic way of identifying cycling markets using the median daily price change. During most days in a cycling market, the change in the average retail price is negative. Falling prices are only occasionally interrupted by one or two days with very large price increases. Therefore, in cycling cities the median of these daily price changes is likely to be negative. In non-cycling cities, prices tend to change slowly in both directions and move only in response to wholesale cost changes. Therefore, over time the median change in retail price in a non-cycling market is likely to be very close to zero as long as wholesale prices are not strongly trending up or down during the period.

The median price change measure does a good job of separating markets that appear to have cycles from those that do not. A large number of cities in our sample have a median price change very close to zero while the rest exhibit a distribution of median price changes that are significantly negative. Though there are a few cities in the sample that show weak cycling characteristics, cities with a $Median\Delta p > -.1$ cents per gallon tend to be clearly non-cycling and cities with $Median\Delta p < -.2$ exhibit very distinct cycles. Therefore, we will study price passthrough in

cycling and non-cycling cities using these two distinct groups.⁴

Estimating Price Passthrough

Non-Cycling Markets

For markets that do not exhibit Edgeworth price cycles we will use an error correction model to estimate the lagged response of retail prices to changes in wholesale costs. Most of the existing literature on gasoline price response uses similar methods.⁵ The main reason for this is that daily (or weekly) wholesale gasoline prices typically resemble a non-stationary process, and in this case, retail and wholesale prices are likely to be cointegrated. Estimating an error correction model in the spirit of Engle and Granger (1987) produces consistent estimates in the presence of cointegration. The resulting coefficient estimates are used to construct a response function that describes how a hypothetical cost change is passed through to the retail price.

The empirical model is adapted from the following error correction model of retail prices, p , and wholesale costs, c , for period t in market m :

$$\Delta p_{mt} = \sum_{i=0}^{I-1} \beta_i \Delta c_{m,t-i} + \sum_{j=1}^{J-1} \gamma_j \Delta p_{m,t-j} + \theta z_{mt} + \epsilon_{mt}. \quad (1)$$

$$z_{mt} = p_{m,t-1} - \left(\phi c_{m,t-1} + \sum_{m=1}^M (v_m \text{MARKET}_m) \right) \quad (2)$$

where:

$$\Delta p_{mt} = p_{mt} - p_{m,t-1} \quad \text{and} \quad \Delta c_{mt} = c_{mt} - c_{m,t-1}$$

$$\text{MARKET}_m = \text{City fixed effects}$$

The z_{mt} term represents how far the retail price in the previous period, $p_{m,t-1}$, was from its typical long run equilibrium level as determined by the wholesale cost level, $c_{m,t-1}$, and the average profit margin, MARKET_m , for that particular market. The coefficient θ represents the speed with which prices tend to revert back to this long run equilibria. The previous literature concludes that a more flexible model is needed to account for the asymmetric adjustment patterns observed in gasoline

⁴Based on these classifications we use 39 cities in our non-cycling sample and 33 cities in our cycling sample.

⁵See, for example, Bachmeier and Griffin (2003) and Lewis (2005).

prices. As is typical, we estimate the lagged cost change and price change coefficients (β 's and γ 's) separately for cost increases and decreases. In addition, we allow the error correction coefficient (θ) to take on different values when price is above or below its long run equilibrium relationship (ie. when z_{mt} is positive or negative). This relaxed functional form is commonly referred to as a threshold error correction model.

The number of lagged cost and price changes included the estimation must also be specified. Unfortunately, statistical procedures to determine the proper lag length do not work well in our application. Previous studies that use similar models tend to include lags of cost changes going back one to two months.⁶ However, there are frequently isolated lags beyond 2 months that are statistically significant but have an economically insignificant impact on predicted price response paths. Therefore, we limit our lags to 40 days of cost change and 15 days of price changes. Alternative specifications with different lag lengths yield very similar response estimates. Huber-White robust standard errors calculated to account for possible heteroskedasticity.

The following model is estimated:

$$\Delta p_{mt} = \sum_{i=0}^{39} (\beta_i^+ \Delta c_{m,t-i}^+ + \beta_i^- \Delta c_{m,t-i}^-) + \sum_{j=1}^{15} (\gamma_j^+ \Delta p_{m,t-j}^+ + \gamma_j^- \Delta p_{m,t-j}^-) + \theta^+ z_{mt}^+ + \theta^- z_{mt}^- + \epsilon_{mt}. \quad (3)$$

$$z_{mt} = p_{m,t-1} - \left(\phi c_{m,t-1} + \sum_{m=1}^M (v_m \text{MARKET}_m) \right). \quad (4)$$

Following Engle and Granger (1987) we use a two-step estimation procedure. The first step uses Equation 4 to recover an estimate of z_{mt} equal to the residual of an OLS regression of the retail price on the wholesale price and market fixed effects. The second step then estimates Equation 3 by using the residual from the first stage, $\hat{z}_{mt}$, in place of z_{mt} .

The estimated coefficients are not easily interpretable, especially given the nonlinearities involved in the model. What is more interpretable and more relevant is the cumulative response functions of retail prices to wholesale cost shocks generated from these coefficients. Following the standard practice in the literature, we construct cumulative response functions that trace out

⁶Borenstein et al. (1997) use bi-weekly data and include three lags. Lewis (2005) uses weekly data and includes 7 cost lags and 4 price lags.

Figure 2: Estimated Retail Price Response in Non-Cycling Markets.

the impact of positive and negative cost shocks on final retail prices taking into account all the nonlinearities and all the transmission channels through which the shocks operate. The predicted effect on price n periods after a cost change includes the direct effect of the past cost change (β_{t-n}) plus the indirect effects from the resulting price changes in the previous periods $n-1$ periods (γ_j 's), and the error correction effect.⁷

Figure 2 shows the typical speed of response to a cost increase and a cost decrease. Not surprisingly, the response to negative cost shocks is much slower. A large fraction of cost increases are passed through quickly, with prices responding to almost 75% of the cost change after one week. However, prices are slow to fully respond to cost changes, even for positive shocks. It takes three weeks following a cost increase for prices to approach full passthrough. For negative cost shocks it takes three weeks for prices to incorporate 75% of the cost change and nearly six weeks to approach full passthrough. This slow and asymmetric response is consistent with the findings of previous studies of gasoline pricing. Surprisingly, the next section shows that prices respond much more quickly in cycling markets than in the non-cycling markets described above.

⁷For a more detailed discussion of the construction of cumulative response functions see the Appendix of Borenstein et al. (1997).

Cycling Markets

Unlike noncycling markets, it is common for a station in a cycling market to change its price almost every day, either to undercut a competitor or to raise price during the start of a new cycle. The price movements are also asymmetric, with occasional large price increases followed by many days of small price undercuts. The standard error correction model described in the previous section is well suited to study non-cycling markets that exhibit a single steady state price–cost relationship in the long run. However, it is not ideal for studying price responses in cycling markets where prices rise and fall in predictable ways for reasons other than changes in cost. Building a model specific to cycling markets that accounts for the two phases of the cycle allows a much crisper estimation of price movements and the speed of response of cost changes. This is done with a latent regime Markov switching regression framework.

Two cycle phases are suggested by both the theory of Edgeworth cycles and by the data:

1. the relenting phase (phase “R”), and
2. the undercutting phase (phase “U”)

with discrete switching between the two. The nature of Edgeworth cycles is that the phases are correlated over time. Undercutting phases tend to persist for many consecutive days as firms undercut one another. Relenting phases tend to last just a few days as the firms in the market respond to a large price increase by one firm by increasing prices themselves. Therefore, the current regime carries information about the likelihood of the phase in the following period. This is well handled by a Markov switching regression framework with two distinct phases and discrete switching. A regular switching model does not have this memory feature.⁸

For a given market, price movements within phase k of the cycle are modeled by the function

$$\Delta p_{mt} = \sum_{i=0}^N \beta_i^k \Delta c_{m,t-i} + \alpha^k p_{m,t-1} - \gamma^k c_{m,t} + \sum_{m=1}^M (\tau_m^k \text{MARKET}_m) + \varepsilon_{mt}^k, \quad k = R, U \quad (5)$$

$$\Delta p_{mt} = p_{mt} - p_{m,t-1} \quad \text{and} \quad \Delta c_{mt} = c_{mt} - c_{m,t-1}$$

⁸For further discussion of the advantages of the latent regime Markov switching regression framework for modeling Edgeworth cycle price movements see Noel (2007a) & Noel (2008).

for the relenting ($k = "R"$) and undercutting ($k = "U"$) phases.⁹ The probability that the market switches from regime k to the relenting phase "R" in a given period is

$$\begin{aligned}\lambda_{mt}^{kR} &= \Pr(I_{mt} = "R" \mid I_{m,t-1} = k, W_{st}^k) \\ &= \frac{\exp(W_{mt}^k \psi^k)}{1 + \exp(W_{mt}^k \psi^k)}, \quad k = R, U\end{aligned}\tag{6}$$

$$W_{mt}^k \psi^k = \sum_{i=0}^N \theta_i^k \Delta c_{m,t-i} + \zeta^k p_{m,t-1} - \omega^k c_{m,t}\tag{7}$$

where $\lambda_{mt}^{kU} = 1 - \lambda_{mt}^{kR}$ to satisfy the adding up constraint. The error terms ϵ_{mt}^k are normally distributed with mean zero and variance σ_i^2 . The phase in which prices are found to fall gradually each day is called the undercutting phase, and the other is called the relenting phase. The core model parameters ($\beta^k, \alpha^k, \gamma^k, \theta^k, \zeta^k, \omega^k, \tau^k, \kappa^k$) are simultaneously estimated by the method of maximum likelihood and Newey-West standard errors are reported to be conservative.

Intuitively, the largest impact of a cost change occurring during the undercutting phase in a cycling market is that the next relenting phase will come sooner following a cost increase or will be delayed following a cost decrease. Similarly, if costs change during a relenting phase, an increase in costs will cause prices to rise higher before a new undercutting phase begins, whereas a decrease in costs will cause prices not to rise as far before a new undercutting phase begins. In the empirical model this is reflected by the importance of the last period retail price $p_{m,t-1}$ and the current rack price c_{mt} .

Each of c_{mt} and $p_{m,t-1}$, which are known to firms when they decide current prices determining Δp_{mt} , appear in both Equations 5 & 7 for each regime k . The difference between them, $p_{m,t-1} - c_{mt}$, represents the market's current relative *position* in the cycle. The theories of Edgeworth cycles predicts that current relative position impacts the probability of a switch between phases and potentially the magnitude of the price changes within each phase.¹⁰ When firms undercut one another down toward the bottom of the cycle, so that position approaches zero, it should be increasingly likely that a new relenting phase will begin ($\zeta^U < 0, \omega^U > 0$). If in a relenting phase already, that phase should be more likely to continue as long as position is still relatively

⁹While some parameters share symbols with those from the non-cycling market model, they are not intended to represent the same parameter.

¹⁰See Maskin and Tirole (1988) and Noel (2008)

low ($\zeta^R < 0$, $\omega^R > 0$). As noted, the cycles are highly predictable and they tend to have a roughly constant amplitude for a given market, meaning that firms are reinstating a relatively constant markup at the top of the cycle each time.

The current position in the cycle also affects the magnitudes of price changes conditional on phase. When undercutting firms approach the bottom of the cycle, the undercuts, if they still occur, should become smaller in absolute value as firms enter the war of attrition ($\alpha^U < 0$, $\gamma^U > 0$, since α and γ are negative numbers). The closer cycle position is to zero, the greater should be the expected price change in the relenting phase ($\alpha^R < 0$, $\gamma^R > 0$). This latter relationship is mitigated some by the market average nature of the price data and the fact that relenting phases take two days to complete marketwide. However, this has no appreciable impact on the categorization of phases or parameters of interest. The results will show the effect of current market position is strong, and in each way consistent with the theory.

It is possible that firms could incorporate past cost shocks into current prices with a lag, independent of the current position in the cycle. To account for this, lagged cost changes, $\Delta c_{m,t-i}$, are included in the Equations 5 & 7 for each regime k . This allows lagged cost changes to affect both the magnitude of price changes conditional on phase, for each phase, and also the switching probabilities out of each phase. A testing down approach was taken to arrive at twenty lags. The small standard errors in the model estimation, itself due to the high predictability of pricing along the cycle, result in occasional statistical significance on one of the longer lagged difference terms, but the results are effectively unchanged even when the number of lags included is reduced to as low as six or seven.

The specification also includes a full set of city dummies in each phase equation to allow pricing behavior and cycle characteristics to vary across cities. In addition, the wholesale and retail price series are demeaned to remove any differences in average markups across cities.¹¹

The parameter estimates on current market position all yield signs consistent with the theories of Edgeworth cycles. Table 1 shows the parameter estimates from the main model on the four $p_{m,t-1}$ and c_{mt} terms, from the within-phase Equation 5 and the switching probabilities

¹¹Results are very similar when city dummies are not included in the price change equations or when the wholesale and retail series are not demeaned. This is in large part because of the uniformity of cycle patterns across cycling cities.

Table 1: Cycling Model Coefficients on $p_{m,t-1}$ and $c_{m,t}$ from Equations 5 & 7.

<u>Variable</u>	<u>Coefficient Estimates</u>	
Price Equation:	U	R
$c_{m,t}$	$\gamma^U : .0470^*$ (.0020)	$\gamma^R : .1762^*$ (.0707)
$p_{m,t-1}$	$\alpha^U : -.0502^*$ (.0020)	$\alpha^R : -.2161^*$ (.0661)
Switching Equation:	U $\rightarrow$ R	R $\rightarrow$ R
$c_{m,t}$	$\omega^U : .2539^*$ (.0160)	$\omega^R : .0073$ (.0169)
$p_{m,t-1}$	$\zeta^U : -.2486^*$ (.0158)	$\zeta^R : -.0008$ (.0169)

*Indicates statistical significance at the 1% level.

Equation 7, for each phase. As expected, firms are more likely to enter a relenting phase when undercutting closer to the bottom of the cycle ($\zeta^U = -.249 < 0$, $\omega^U = .254 > 0$), and begin undercutting when near the top (i.e. stop relenting near the top) ($\zeta^R = -.001 < 0$, $\omega^R = .007 > 0$). Relenting phase price increases are higher and undercutting phase price decreases are lower in absolute value the closer firms are to the bottom of the cycle. ($\alpha^R = -.216 < 0$, $\gamma^R = .176 > 0$, $\alpha^U = -.050 < 0$, $\gamma^U = .047 > 0$).

Overall, there are 80 lagged rack price change coefficients in the model: twenty lagged cost coefficients in the price change equation for each of the two phases of the cycle, and twenty more in the switching probability equation for each phase. As in the estimation for non-cycling cities, the individual coefficients themselves are not particularly meaningful given the nonlinearities in the model. So we instead focus our discussion on the estimated cumulative response functions generated from these coefficients.

Unlike in the non-cycling cities, how the price responds to a cost change in a cycling market depends on the position of the price within the cycle when the cost change occurred. Therefore, we perform the following simulations to generate an average price response path over many different starting price positions. Fifteen hundred retail price paths are simulated following

a positive, permanent shock to the rack price and fifteen hundred following a negative, permanent shock. Consider a shock Δc . For each simulation, a new draw is taken of the core parameter vector $(\beta^k, \alpha^k, \gamma^k, \theta^k, \zeta^k, \omega^k, \tau^k, \kappa^k)$ from its distribution. The simulations are carried out twenty periods after the shock to run as long as the longest lag in the estimation. Equal numbers of simulations are performed across a range of fifteen starting values of p_{mt} that span the possible values of p_{mt} along the path of the cycle.¹² Market dummies are set equal to $1/M$, where M is the number of cycling markets.¹³ Mean cumulative response functions for increases and decreases are presented along with their respective standard errors calculated from the variation across simulations.

In cycling markets the speed of response will be different for different sizes of cost shocks, because the impact of the shock on the current position within the cycle depends on the relative size of the shock and the amplitude of the cycle. Therefore, the cumulative response functions are reported for two different magnitudes of cost changes. The first value of Δc is mean daily non-zero cost change, ($\Delta c = 2.69$), and second is the mean five-day change ($\Delta c = 4.92$). The first magnitude reflects a typical one day cost change, while the latter is intended to represent larger cost movements that are common in the industry but are often spread out over a period of days.¹⁴

The speed of responses of retail prices to cost shocks in cycling markets are shown in Figures 3 and 4. The responses to both increases and decreases are much faster than in non-cycling markets. In Figure 3, a cost increase of 2.69 cents is fully passed through to the retail price in just five periods (five days). For an equal sized cost decrease, one hundred percent passthrough is achieved in just seven days. In Figure 4 using $\Delta c = 4.92$, the mean of the five day rack price change, one hundred percent passthrough is achieved five days after an increase and nine days after a decrease. Compared to non-cycling cities, these response times are exceptionally short, for

¹²Computer run time is exponential in the number of periods in the simulation, and linear in the number of simulations and number of p_{mt} points used in the average. Running the simulation out twenty periods, with three thousand simulations (2 x 100 simulations per p_{mt} point x 15 p_{mt} points) takes four days on 2GB of RAM. Reducing the number of periods to fifteen, it takes an hour. Passthrough is largely complete in ten or twelve periods, so little is lost from the reduced number of periods. Using ten periods and three hundred thousand simulations (2 x 3000 simulations per p_{mt} point x 50 p_{mt} points) yields the same conclusions, as do other feasible combinations.

¹³Results do not meaningfully change if all market dummies are set to zero (effectively choosing the estimates of the omitted city), or if any one particular market is chosen and its dummy is set to one. The overall conclusions hold firm.

¹⁴Simulations were conducted across a range of other values of Δc , from 0.0001 up to 8. The results confirm the patterns of responses and asymmetry discussed below.)

Figure 3: Estimated Retail Price Response in Cycling Markets (Cost Change = 2.69 cents).

Figure 4: Estimated Retail Price Response in Cycling Markets (Cost Change = 4.92 cents).

both cost increases and decreases.

It is the mechanism of the price cycle itself allows cost changes to be passed through more quickly in these cities. The amplitude of a cycle is roughly constant in a given market, so a cost shock effectively moves future peak prices and future trough prices up or down by a roughly equal amount relative to earlier peak and trough prices. A negative shock lowers them, a positive shock raises them. After a cost increase, undercutting firms will reach the next trough a bit sooner than otherwise, and when they do they will relent to a new higher price and form the next peak. This incorporates the cost increase at once, relative to the last peak. However, because firms were in between the peak and trough price at the time of shock, the amount passthrough can be well over 100% at that moment. By the time firms return to the same relative position in the cycle (e.g. half way between top and bottom, if that is when the shock originally occurred), passthrough of cost increases is largely complete around 100%. Conversely, a cost decrease lowers the new trough price and gives firms additional room to undercut along the downward portion of the cycle. They undercut through the old trough price to the new bottom of the cycle. By the time they return to the same relative position in the cycle as when the shock occurred, the cost shock has effectively fully been incorporated into the price. Through this mechanism, cycles eliminate the frictions in price movements. Because cycle troughs and peaks are largely determined by the current cost (and price) level, cycling markets naturally absorb cost shocks within the period of one cycle.

As in non-cycling markets, the response to cost increases is faster than the response to cost decreases. The difference is significant for the first ten periods in both Figures 3 and 4. However, in markets with cycles this asymmetry in response is generated by the cycling mechanism itself. Cost increases often trigger a new relenting phase and relenting phase price changes are very large. Cost decreases are always followed by additional undercuts, and undercuts tend to be small. As a result, the cycle generates some asymmetry in response during the first cycle following a cost shock. For larger cost shocks this asymmetry is greater because large cost increases are more likely to trigger a new relenting phase, and undercutting is likely to persist longer after large cost decreases. Nevertheless, the asymmetry disappears once one full cycle is complete and the cost changes have been fully passed through.

Discussion

The most striking result from our comparison of cycling and non-cycling markets is that cost changes are passed through much more quickly in markets with price cycles. It takes at least three times as long for prices in non-cycling cities to fully reflect a change in costs. The finding strongly suggests that the existence of cycles significantly affects the passthrough of cost shocks. Moreover, the continuous movement of prices within the cyclical equilibrium provides a clear mechanism revealing how costs could be passed on more quickly.

Explanations for delayed price response in non-cycling markets often involve either focal prices or menu costs. Some theories suggest that past price levels may be used as a focal price, either by tacitly colluding firms (Borenstein et al. (1997)) or by imperfectly informed consumers (Lewis (2005)), causing prices to be sticky. In a cycling market, there is little potential for focal prices to develop as prices fluctuate continuously along the cycle. Similarly, the theory that firms face menu costs and are reluctant to change prices in response to each cost shock does not make sense in cycling markets. Firms in these markets change prices nearly every day, even in the absence of a cost shock. Rather than prices being sticky, the constant movement of prices makes it easier for firms to incorporate changes in costs.

Although cyclical price movements provide a natural explanation for faster price response, it is possible that other factors are involved. The most likely alternative is that differences in the market structure of competitors or in the demographics of consumers across markets cause cycles to occur and also cause costs to be passed through more quickly. Noel (2007a), Lewis (Forthcoming), and Doyle et al. (2008) all find evidence that the existence of cycles is correlated with various measures of station density and concentration. Using the same data analyzed in this study, Lewis (Forthcoming) identifies that the presence of cycles is most strongly associated with population per station, the share of independent stations, and the ownership concentration (or Herfindal Index) of independent stations.¹⁵ These characteristics of market structure could

¹⁵Independent stations are those that do not advertise or sell the “branded” gasoline of a major refining company such as BP, Exxon, or Shell. Independent stations can be either single station operations or chains of gas stations or convenience stores. Some larger independent chains have recognizable brand names of their own based on their convenience store operations (such as Speedway, Quik Trip, and WaWa), but they do not market the gasoline of a

also impact the speed of price response. Rather than cycles generating faster passthrough, it is possible that the speed of passthrough depends more on market structure than on the existence of cycles. In this case, one might expect to see roughly similar rates of price passthrough in cities with similar market structure characteristics even if one of the cities had cycles and one did not. However, our earlier findings suggest that this is not the case. Markets with and without cycles have dramatically different rates of passthrough despite the fact that many cycling and non-cycling markets have similar characteristics.

One way to examine this more carefully is to relate the speed of price response in each city with its market characteristics (including the presence of cycles). Our previous results are presented as average response speeds for cycling and non-cycling markets, but both models also allow for the estimation of city-specific response speeds. For non-cycling markets, the model in Equation 1 can be altered to estimate separate values of the error correction coefficient, θ , for each market. This allows the speed with which prices revert to their long run equilibrium margin to vary across cities.¹⁶ Response function estimates can then be constructed for each city just as they were for the entire group in Figure 2. The model for cycling markets already includes city fixed effects in the price change equations for each phase (Equation 5). Figures 3 & 4 present the estimated price response based on the average market, but response estimates for each market can be constructed similarly.

We estimate price response functions for each market and compare them by examining the share of a cost change that is passed through after a specific number of days. As an example, consider the share of a cost decrease that has been passed on by the fifth day after the shock. Figure 5 displays four scatterplots illustrating how passthrough in the fifth day relates to the presence of cycles (Panel A) and to the three market structure characteristics that have been found to predict the presence of price cycles in Lewis (Forthcoming).¹⁷ The population per station within

specific refining company.

¹⁶ Another approach would be to estimate the entire error correction model separately for each market in order to construct market specific response function estimates. Given the large number of coefficients on lagged cost and price changes and the relatively short sample period, this results in very imprecisely estimated coefficients and response functions. As a result we effectively restrict coefficients on lagged cost and price changes to be the same across cities while allowing the error correction coefficient to vary.

¹⁷ Figure 5 displays the passthrough of a negative cost change. Graphs showing the passthrough of a positive cost

Figure 5: How Speed of Passthrough Relates to Market Characteristics.

the Metropolitan Statistical Area (Panel B) is constructed using population and retail gasoline establishments data from the 2000 Census and the 2002 Economic Census. The market share of independent stations (Panel C) and the HHI among independent stations (Panel D) are constructed from brand level market shares provided by OPIS.¹⁸

Figure 5 (Panels B-D) reveals that population per station, the market share of independent stations, and the HHI within independent stations all appear to have (at best) a weak positive relationship with the speed of price response. Many cities with very similar values of these market structure characteristics have drastically different passthrough rates. In contrast, the presence of cycles is strongly positively related with the speed of passthrough (Panel A). The distribution of passthrough rates for the cycling cities is completely distinct from that of non-cycling cities, with

change look almost identical except that the observations are shifted up slightly because price increases are passed through more quickly than decreases.

¹⁸For details on the construction of all three variables see Lewis (Forthcoming).

Table 2: Regression of Speed of Price Response on Market Characteristics.

Dependant Variable = Share of Δc Passed Through After 5 Days		
Size of Cost Change:	$ \Delta c = 2.69$	$ \Delta c = 4.92$
Price Cycle Present	.399*	.388*
	(.026)	(.022)
Population per Station	.004	.006
	(.023)	(.017)
Marketshare of Independent Stations	.119	.113
	(.081)	(.072)
HHI within Independent Stations	.246 ⁺	.240*
	(.099)	(.087)
Cost Increase	.450*	.491*
	(.059)	(.047)
Cost Decrease	.132 ⁺	.101 ⁺
	(.059)	(.047)
Adjusted R ²	.960	.970
Number of Observations	138	138

Notes: * and + indicate statistical significance at the 1% and 5% levels respectively. “Cost Increase” and “Cost Decrease” are indicator variables whose coefficients reflect the constant term in the regression when the cost change is positive or negative respectively.

every cycling city having more complete passthrough after 5 days than *any* non-cycling city.

The results of a simple regression confirm that the existence of cycles is the best predictor of the speed of price response. Table 2 reports estimates from regressions of the fraction of a cost change passed through by the fifth day on the market structure variables (for two different magnitudes of Δc).¹⁹ The coefficient on the presence of cycles is by far the strongest and most statistically significant. It suggests that the average share of a cost change passed through in a cycling city is nearly 40 percentage points higher than in a non-cycling city with similar market structure. This cycle indicator alone explains nearly half of all observed variation in the speed of passthrough across cities.

The coefficient on the HHI within independent stations is also statistically significant, but

¹⁹While we present results based on passthrough by the fifth day following a cost change, the regression results using passthrough rates on other days are very similar to those reported here.

it predicts much less of the overall variation in the speed of passthrough. A two standard deviation increase in HHI within independents is associated with a 6 percentage point increase in the share of cost change passed through by the fifth day. Cyclical price movements themselves appear to have a greater direct influence on the speed of price response than the underlying market structure characteristics that make cycles more likely to occur.

Conclusion

Our empirical analysis reveals that prices respond much more quickly to cost shocks in markets with Edgeworth price cycles. The environment of constantly changing prices characteristic of an Edgeworth cycle equilibrium enables cost changes to be fully passed through in one third the time it takes in a typical non-cycling market. We confirm that the existence of cycles generates faster passthrough by studying how it relates to the speed of price response in each city after controlling for various measures of market structure. Prices in cities with cycles respond much more quickly even when compared to cities with similar market characteristics. We interpret these findings as further evidence that prices respond more rapidly to cost shocks as a result of the presence of the price cycle rather than as a result of other market characteristics that are correlated with the presence of price cycles. They confirm the intuition that the constant price movement generated by the price cycle more quickly incorporates changes in cost.

We believe this is one of the first concrete results identifying specifically how the existence of Edgeworth price cycles affects the performance and overall competitiveness of a market. It is also one of the first papers within the sticky price literature to clearly identify a particular market characteristic responsible for such large differences in the speed with which prices respond to cost changes. The inefficiencies generated by slow price passthrough are largely eliminated when cycles are present.

References

- Bachmeier, Lance J. and James M. Griffin**, “New Evidence On Asymmetric Gasoline Price Responses,” *The Review of Economics and Statistics*, August 2003, 85 (3), 772–776.
- Borenstein, Severin, Colin Cameron, and Richard Gilbert**, “Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?,” *The Quarterly Journal of Economics*, February 1997, 112, 306–339.
- Deltas, George**, “Retail Gasoline Price Dynamics and Local Market Power,” *Journal of Industrial Economics*, 2008, 56 (3).
- Doyle, Joeseeph, Erich Muehlegger, and Krislert Samphantharak**, “Edgeworth Cycles Revisited,” June 2008. Working Paper, MIT.
- Eckert, Andrew**, “Retail Price Cycles and the Presence of Small Firms,” *International Journal of Industrial Organization*, February 2003, 21 (2), 151–170.
- Engle, Robert F. and Clive W. J. Granger**, “Co-Integration and Error Correction: Representation, Estimation, and Testing,” *Econometrica*, 1987, 55 (2), 251–276.
- Lewis, Matthew**, “Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market,” July 2005. Working paper. The Ohio State University, Columbus, OH.
- , “Temporary Wholesale Gasoline Price Spikes have Long-lasting Retail Effects: The Aftermath of Hurricane Rita,” *The Journal of Law & Economics*, Forthcoming, 52.
- Maskin, Eric and Jean Tirole**, “A Theory of Dynamic Oligopoly II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles,” *Econometrica*, 1988, 56, 571–599.
- Noel, Michael**, “Edgeworth Price Cycles: Cost Based Pricing and Sticky Pricing in Retail Gasoline Markets,” *Review of Economics and Statistics*, 2007, 89 (2), 324–334.
- , “Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market,” *Journal of Industrial Economics*, 2007, 55 (1), 69–92.

—, “Edgeworth Cycles and Focal Prices: Computational Dynamic Markov Equilibria,” *Journal of Economics and Management Strategy*, 2008, 17 (2), 345–377.

Verlinda, Jeremy A., “Do Rockets Rise Faster and Feathers Fall Slower in an Atmosphere of Local Market Power? Evidence from the Retail Gasoline Market,” *Journal of Industrial Economics*, 2008, 56 (3).

Wang, Zhongmin, “Collusive Communication and Pricing Coordination in a Retail Gasoline Market,” *Review of Industrial Organization*, February 2008, 32 (1), 35–52.