

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†

October 2, 2010

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 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

I. 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 pass-through have important implications for macroeconomic policy. On a more micro level, studying pass-through 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 (Forthcoming) 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 speed gasoline price response. However, overwhelming empirical evidence shows that the process is slow, often taking a month or two to fully incorporate cost shocks into the retail price.

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 (2009); Doyle et al. (2010)], 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 (2009); Doyle et al. (2010)], but the impact that cycles have on the competitiveness of the local market is still not well understood.

In this article 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 in 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 pass-through.

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 pass-through 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.

II. Edgeworth Price Cycles in Retail Gasoline Markets

The Maskin and Tirole (1988) model contains two identical homogeneous product firms competing in an alternating move pricing game. They show that one equilibrium of the game results in a cyclical pattern of prices over time known as an Edgeworth cycle. 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 firms play a war of attrition. Each would like to see prices return back to high levels, but they would prefer the other firm to raise price first. On each turn, a firm mixes between a price equal to marginal cost and price equal to a high, possibly supra-monopoly price. Ultimately one firm raises its price, and in the next period the other follows immediately, just undercutting the price of the first. Then the cycle of undercutting begins again.

Although this model is highly stylized, it turns out to be a very good representation of price competition in certain retail gasoline markets. The predicted pattern of prices strongly resembles the cyclical gasoline price movements observed in these cities. The time series appears as a sawtooth pattern, with consecutive periods of undercutting interrupted by large, occasional price jumps closely synchronized across firms to restore prices back to a profitable level. The alternating move nature of the Maskin and Tirole game also resembles the back-and-forth strategic price jockeying that is often associated with competition between neighboring gas stations. Moreover, Noel (2008) extends the basic Maskin and Tirole model and shows that Edgeworth cycle equilibria exist in a number of other settings, including cases with differentiated firms and with more than two firms.

The theory also provides a clear prediction for how prices will respond to cost changes. In an Edgeworth cycle equilibrium the current profit margin determines when firms decide to restore prices to more profitable levels, and firms restore prices to a specified markup over the current cost. Therefore, the occurrence of a cyclical price restoration fully incorporates any past cost changes into the current retail price. As a result, cost changes in cycling markets should be passed through to prices within the trough-to-trough period of a single price cycle.¹ The exact speed of pass-through depends on when the cost shock occurs within the cycle. If a cost shock happens right before a restoration is to occur then it will be passed on even more quickly. Since cycles in these markets typically last only one to two weeks, this suggests that cost changes should be

¹Maskin and Tirole (1988) consider constant costs but show that restorations only occur once prices reach the level of cost. As cost changes, the level of prices required for non-positive markups that trigger a restoration will change as well. Noel (2008) addresses the impact of cost changes directly and again shows that retail prices return to a minimum markup over cost just prior to each restoration.

passed through in roughly one third the time that is commonly observed in non-cycling markets. Our empirical analysis tests this prediction by estimating pass-through speeds in both cycling and non-cycling markets.

III. Data

In order to compare across market types, our empirical analysis utilizes a broad panel of daily retail and wholesale gasoline prices from 90 cities between January 1, 2004 and August 28, 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.

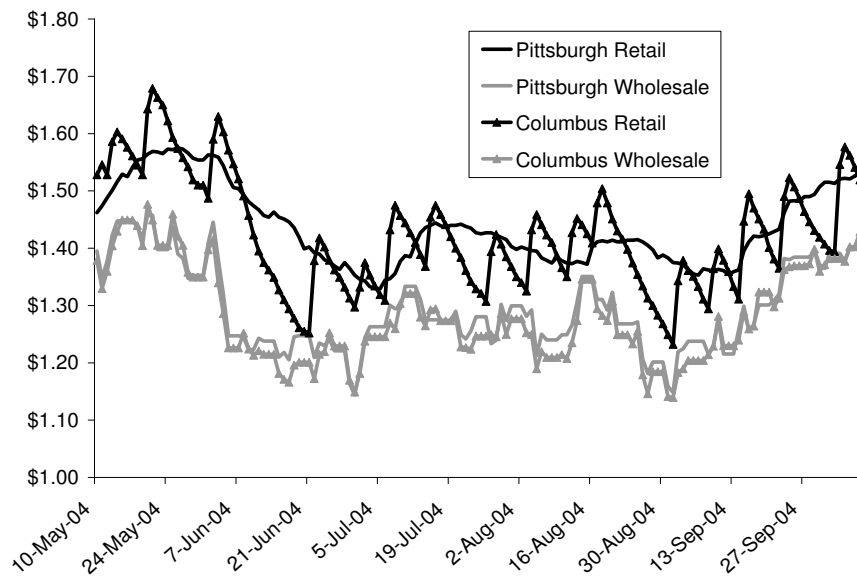
A. *Identifying Markets with Retail Price Cycles*

Analyzing differences in wholesale cost pass-through 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 independent of wholesale price movements. In contrast, Pittsburgh's retail price shows much less day to

²Retail prices are for a gallon of regular (87 octane) gasoline with all relevant taxes removed. According to OPIS, price information comes from a number of sources including fleet credit card transactions and direct feeds from stations. Prices are observed from all types of stations including refinery owned stations, lessee dealers, and independent retail chains.

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.



day fluctuation and most of the movement in the retail price appears to correspond with changes in the wholesale price.

Lewis (2009) 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. There is no evidence that any of the cities in the sample switch to or away from cycles during the period.⁴ Though there are a few cities in the sample that show weak

⁴In his study of Canadian gasoline markets Noel (2007a) also considers a third type of equilibrium, a sticky

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 pass-through in cycling and non-cycling cities using the 72 cities that fall into these two distinct groups.⁵

IV. Estimating Price Pass-through

A. 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.⁶ 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 :

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

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

where:

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

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

or rigid price regime, in which retail prices remain unchanged for long periods and do not respond to small cost fluctuations. Visual inspection reveals that there are clearly no cities in our U.S. sample that would classify in this category, so we simply group cities as cycling or non-cycling.

⁵Based on these classifications there are 39 cities in our non-cycling sample and 33 cities in our cycling sample.

⁶See, for example, Bachmeier and Griffin (2003) and Lewis (Forthcoming).

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 price tends to revert back to this long run equilibrium. The previous literature concludes that a more flexible model is needed to account for the asymmetric adjustment patterns observed in gasoline 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.⁷ 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:

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

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

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

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

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} .

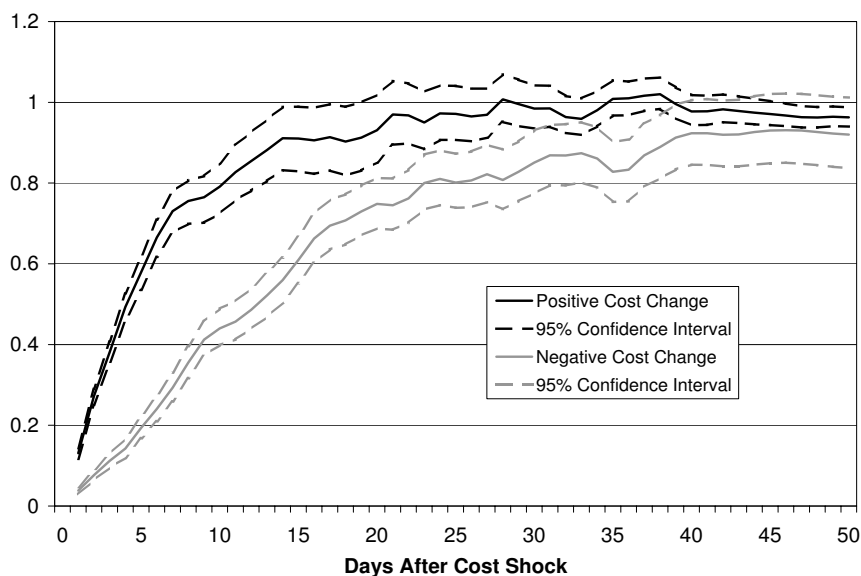
Coefficient estimates for Equation 3 are reported in Appendix Table A1. The estimated coefficients are not easily interpretable, especially given the nonlinearities involved in the model. What are more interpretable and more relevant are the cumulative response functions of retail prices to wholesale cost shocks generated from these coefficients. Following the standard practice in the literature, we construct cumulative response functions that trace out 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 $n-1$ periods (γ_j 's), and the error correction effect.⁸ Standard errors for the points on the cumulative response function are estimated using a bootstrap procedure based on the estimated covariance matrix of the coefficients.⁹

Figure 2 shows the typical speed of response to a cost increase and a cost decrease. Note that the absolute value of the cumulative response function for a cost decrease is shown so both cumulative response functions (for increases and decreases) can be shown on the same axis. 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 pass-through. For negative cost shocks it takes three weeks for prices to incorporate 75% of the cost change and nearly six weeks to approach full pass-through. 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.

⁸For a more detailed discussion of the construction of cumulative response functions see the Appendix of Borenstein et al. (1997).

⁹One thousand different sets of coefficients (for each parameter in Equation 3) are drawn from a multivariate normal distribution with means equal to the coefficient estimates and a covariance matrix equal to the estimated covariance matrix. Cumulative response functions are constructed for each set of coefficient draws and the standard deviation for each point on the response function across the 1000 simulations is used as an estimate of the standard error.

FIGURE 2: ESTIMATED RETAIL PRICE RESPONSE IN NON-CYCLING MARKETS.



B. *Cycling Markets*

Unlike non-cycling 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. In these markets prices rise and fall in predictable ways for reasons other than changes in cost. Furthermore, the impact of a cost change on price will vary dramatically depending the current position of price within the cycle. The error correction model does not allow for the types of price movements. As a result, estimates from the error correction model will be a very poor predictor of prices in a cycling markets, and may even introduce some unknown bias into the price response estimates.

Our goal is to compare the speed of price response in cycling markets to that in non-cycling markets, so we want to predict price movements as accurately as possible in each market. For the cycling markets we accomplish this by constructing a model that specifically accounts for the

two phases of the cycle and allows the current state of the cycle to impact how prices are passed through. This is done using a latent regime Markov switching regression framework.¹⁰

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.¹¹

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

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

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

for the relenting ($k = “R”$) and undercutting ($k = “U”$) phases.¹² The probability that the market goes from phase k to the relenting phase “R” in a given period is

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

¹⁰Although the error correction model is clearly the wrong choice for estimating price response in cycling markets, we have, as a robustness check, estimated response in these markets using the same model used to estimate non-cycling market price response. We discuss these results briefly at the end of this section.

¹¹For further discussion of the advantages of the latent regime Markov switching regression framework for modeling Edgeworth cycle price movements see Noel (2007a) & Noel (2008).

¹²While some parameters share symbols with those from the non-cycling market model, they are not intended to represent the same parameter.

where

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

and $\lambda_{mt}^{kU} = 1 - \lambda_{mt}^{kR}$ 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 ($\beta^k, \alpha^k, \gamma^k, \theta^k, \zeta^k, \omega^k, \tau^k, \kappa^k$) and their corresponding variance-covariance matrix are simultaneously estimated by the method of maximum likelihood. The likelihood function is for a Markov switching regression and adapted from that in Noel (2007a) and Cosslett and Lee (1985). 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. Then during that relenting phase, prices will rise higher following a cost increase than it otherwise would have, and not as high after a cost decrease than it otherwise would have. 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 phase 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 (Maskin and Tirole (1988), Noel (2008)) make predictions as to how current relative position impacts the probability of switching between phases and the magnitudes of price changes. In the Maskin and Tirole (1988) model, marginal costs are constant, while in the Noel (2008) model, marginal costs are i.i.d. over time.

Both models show that when firms undercut one another down toward the bottom of the cycle, so that position approaches zero, it becomes more likely that a new relenting phase will begin ($\zeta^U < 0, \omega^U > 0$).¹³ Noel (2008) further shows that in the new relenting phase the magnitude of the price change will be larger when the current position is lower ($\alpha^R < 0, \gamma^R > 0$).¹⁴ This latter

¹³Specifically, in the Maskin and Tirole (1988) model, the probability of relenting goes from zero when price is greater than marginal cost to a positive value less than one when price is equal to marginal cost. In Noel (2008), the probability of relenting increases smoothly at low prices.

¹⁴In the Maskin and Tirole (1988) constant cost model, the top-of-the-cycle and bottom-of-the-cycle prices are

relationship may be mitigated in our results some because we use market average price data and it usually takes two days for relenting phases to complete marketwide.¹⁵

The effect of cycle position on the probability of switching from relenting to undercutting is more subtle. Maskin and Tirole (1988) and Noel (2008) each show that when a firm raises price, it does so to the top of the cycle in a single move. Therefore relenting phases—for a given station—always last one period and are not affected by any variable, including cycle position ($\zeta^R = 0, \omega^R = 0$). However, a marketwide relenting phase lasts more than one period, with some stations relenting earlier and some later. If the process extends more than a few days, the probability of continuing in a marketwide relenting phase may be higher if cycle position happens to still be low ($\zeta^R < 0, \omega^R > 0$). When position is low, there are relatively more stations yet to relent. As a result, it is more likely that the average price will continue to increase the following period.

Lastly, Maskin and Tirole (1988) and Noel (2008) show that as cycle position gets smaller, and firms enter the war of attrition, the magnitude of the price undercut should get smaller as well ($\alpha^U < 0, \gamma^U > 0$, since α and γ are negative numbers).¹⁶

It turns out that the speed with which prices respond to cost shocks, the focus of this article, depends primarily on how current cost shocks affect cycle position and how that affects the new phase of the cycle. Nonetheless, it is possible that firms also incorporate past cost shocks into current prices, independent of current position, and we account for this. We include lagged cost changes, $\Delta c_{m,t-i}$, in the Equations 5 & 7 for each phase 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 eighteen lags. The small standard errors in the model estimation, itself due to the predictability of pricing along the cycle, result in occasional statistical significance on longer lagged difference terms, but the results are effectively unchanged even when the number of lags included is reduced to a few.

constant, so the magnitude of the relenting phase price increase is naturally constant as well.

¹⁵This has no appreciable effect on the categorization of regimes.

¹⁶In Maskin and Tirole (1988) model, the undercut is by one unit on the price grid while price is above marginal cost, and zero when price is equal to marginal cost. In Noel (2008), the size of undercuts along the cycle depends on the market structure considered, but under reasonable market structures with more than two firms, the average change in price is lower in absolute value when cycle position is lower.

TABLE 1: CYCLING MODEL COEFFICIENTS ON $p_{m,t-1}$ AND $c_{m,t}$ FROM THE PRICE EQUATION AND SWITCHING EQUATION.

Variable	Coefficient Estimates	
Price Equation:	Undercutting (U)	Relenting (R)
$c_{m,t}$	γ^U : .0470* (.0020)	γ^R : .1762* (.0707)
$p_{m,t-1}$	α^U : -.0502* (.0020)	α^R : -.2161* (.0661)
Switching Equation:	U \rightarrow R	R \rightarrow R
$c_{m,t}$	ω^U : .2539* (.0160)	ω^R : .0073 (.0169)
$p_{m,t-1}$	ζ^U : -.2486* (.0158)	ζ^R : -.0008 (.0169)

Notes: This table reports the estimates of the coefficients on wholesale price, $c_{m,t}$, and lagged retail price, $p_{m,t-1}$, in the price equation and the switching equation of the model for both the undercutting and relenting phases of the cycle. * indicates statistical significance at the 1% level.

The specification also includes a full set of city dummies in the relenting phase equation and a second full set in the undercutting phase equation. Coefficients are specific not only to the city but to the regime, τ_j^k , to allow regime-specific pricing behavior and cycle characteristics to flexibly vary across cities.¹⁷ In addition, the wholesale and retail price series are demeaned to remove any differences in average markups across cities.¹⁸

Table 1 shows the coefficients on the key cycle position variables $p_{m,t-1}$ and c_m taken from the full model. These two variables appear in the two within-phase Equations 5 and the two switching probabilities Equations 7 and it turns out these are the main drivers of response speed asymmetry discussed below.¹⁹ The complete set of coefficient estimates from the regression, including those on all seventy-two lagged rack price changes, are reported in the appendix.

¹⁷See Noel (2007a) for derivations of measures of cycle characteristics (period, amplitude, asymmetry) as functions of core parameter estimates.

¹⁸Results 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.

¹⁹Results are similar in the simplified model that includes cycle position variables but excludes lagged rack price changes.

The results on $p_{m,t-1}$ and c_{mt} show that, consistent with the theory, 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 that the relenting phase price increases are higher the closer firms are to the bottom. ($\alpha^R = -.216 < 0$, $\gamma^R = .176 > 0$).

The effect of cycle position on the probability of continuing in the relenting phase is insignificantly different from zero, consistent with the theory, though the weak signs are consistent with the effect of data averaging ($\zeta^R = -.001 < 0$, $\omega^R = .007 > 0$). Finally, the size of price undercuts is smaller in absolute value the closer to the bottom of the cycle, ($\alpha^U = -.050 < 0$, $\gamma^U = .047 > 0$), as expected.

The focus of the article is to compare response speeds across cycling and non-cycling cities. As with the non-cycling cities model, looking at the individual coefficients does not give a good sense of the overall response speeds the coefficients imply. So we instead we follow the standard practice in the literature and focus our discussion on the estimated cumulative response functions to cost changes 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 Δc to the rack price and fifteen hundred following a negative, permanent shock $-\Delta c$. 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 eighteen periods. 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.²⁰ Market dummies are set equal to $1/M$, where M is the number of cycling markets.²¹ Each simulation yields a cumulative response

²⁰Computer 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 eighteen periods, with three thousand simulations (2 x 100 simulations per p_{mt} point x 15 p_{mt} points) takes two days on 2GB of RAM. Reducing the number of periods to fifteen, it takes an hour. Pass-through 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.

²¹Results 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.

function for a cost increase and one for a decrease. The mean cumulative response functions across simulations for increases and decreases are then calculated and reported. Standard errors are calculated from the variation in the cumulative responses across simulations for each time t , which in turn is generated by variation in the core parameter estimate draws.

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 the 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.²²

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 pass-through 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 pass-through 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 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 of pass-through can be well over 100% at that moment. By the time firms return to the same relative position in the cycle (e.g.

²²Simulations 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).

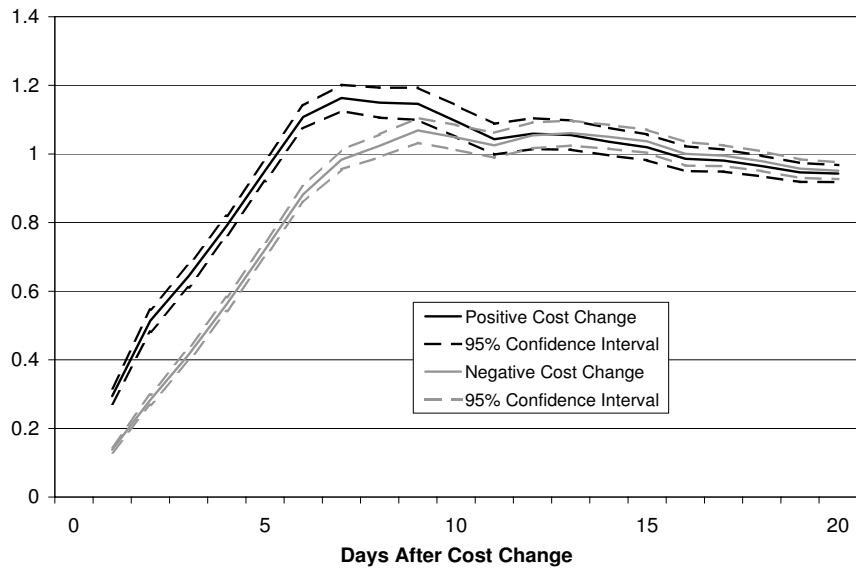



FIGURE 4: ESTIMATED RETAIL PRICE RESPONSE IN CYCLING MARKETS (COST CHANGE = 4.92 CENTS).



half way between top and bottom, if that is when the shock originally occurred), pass-through 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 in large part 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.

The Markov switching model does a very good job of predicting price movements and price response in cycling markets, and it is clearly superior to using a standard error-correction model that is unable to account for cyclical price movements. Nevertheless, for robustness, we have also estimated price response in cycling markets using the same error correction model used for non-cycling markets.²³ Reassuringly, the estimates still suggest that prices respond much more quickly to cost changes in cycling markets than in non-cycling markets. In general the estimates imply that all cost shocks are completely passed through to prices within around 10 days. The confidence intervals around the price response functions are much larger for cycling markets, but this is not surprising given that error correction model is clearly misspecified when cycles are present. Despite this, the results suggest that our findings of faster price response in cycling

²³The results of these additional regressions can be found in the Supplemental Content section of this journal's Web site, available at <http://www.mitpressjournals.org.ezp-prod1.hul.harvard.edu/>.

markets are not simply a product of the Markov model.

V. Discussion

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 (Forthcoming)), causing prices to be sticky. 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 pass-through 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.

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.

The cyclical price mechanism provides a natural explanation for why prices respond faster in cycling markets. However, it is possible that some other market characteristic causes price cycles to occur and also causes prices to adjust more quickly. Therefore, we also attempt to empirically test whether faster price response is generated by price cycles themselves or whether it appears to result from some other aspect of market structure.

The Edgeworth cycle theory suggests that cyclical equilibria may be more likely to occur in markets with some smaller firms that operate fewer stations since these firms have more to gain by undercutting their rivals.²⁴ Similarly, other market characteristics that increase demand cross-elasticities and make undercutting rivals more profitable, such as greater station density, a larger

²⁴See Eckert (2003).

proportion of independent stations, or lower per capita income levels, might increase the likelihood of price cycles. The empirical evidence generally supports these predictions. Noel (2007a), Lewis (2009), and Doyle et al. (2010) all find evidence that cycles are more likely to exist in markets with higher station density and lower levels of market concentration. Using the same data analyzed in this study, Lewis (2009) 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.²⁵

These characteristics of market structure could also impact the speed of price response.²⁶ Rather than cycles generating faster pass-through, it is possible that the speed of pass-through depends more on market structure than on the existence of cycles. In this case, one might expect to see roughly similar rates of price pass-through 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 pass-through 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). If the underlying market characteristics are responsible for faster price response, rather than the cycles themselves, then the presence of cycles should not be a strong indicator of the speed of response once other market factors are controlled for. We test this using a simple cross-sectional regression analysis.

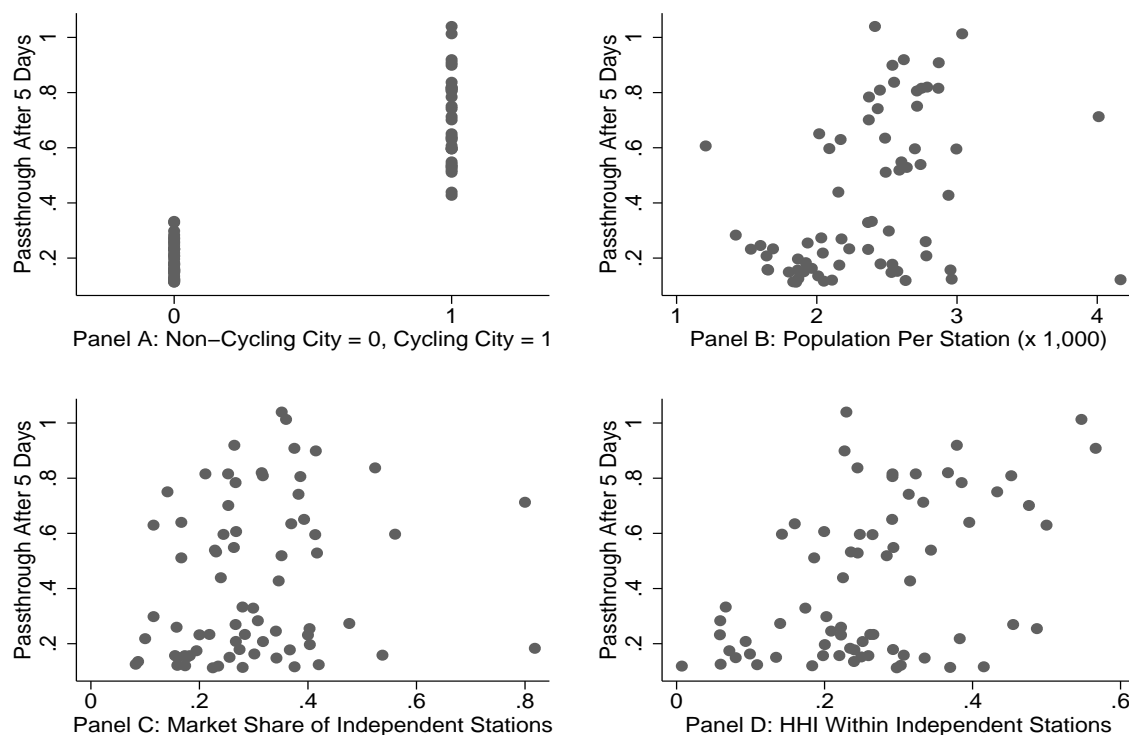
Our previous results are presented as average response speeds for cycling and non-cycling markets, but both models also allow for the estimation of city-specific response speeds. For non-cycling markets, the model in Equation 1 can be altered to estimate separate values of the error correction coefficient, θ , for each market. This allows the speed with which prices revert to their long run equilibrium margin to vary across cities.²⁷ Response function estimates can then

²⁵Independent 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 specific refining company.

²⁶Consumer search behavior and collusion are the two most common explanations of delayed and asymmetric price response in gasoline markets. Greater density and less concentration could reduce consumer’s search costs or make it more difficult for stations to collude, and impact response speed as a result.

²⁷Another approach would be to estimate the entire error correction model separately for each market in order to

FIGURE 5: HOW SPEED OF PASS-THROUGH RELATES TO MARKET CHARACTERISTICS.



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 pass-through 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

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.

the presence of price cycles in Lewis (2009).²⁸ The population per station within 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.²⁹

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 pass-through rates. In contrast, the presence of cycles is strongly positively related with the speed of pass-through (Panel A). The distribution of pass-through rates for the cycling cities is completely distinct from that of non-cycling cities, with *every* cycling city having more complete pass-through 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).³⁰ Columns 3 & 4 include interactions between the presence of cycles and the market characteristics. In all specifications, market characteristics are demeaned to ensure that the coefficient on the presence of cycles variable reflects the difference in the speed of price response between cycling and non-cycling cities with average market characteristics.

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 pass-through across cities.

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

²⁸Figure 5 displays the pass-through of a negative cost change. Graphs showing the pass-through of a positive cost change look almost identical except that the observations are shifted up slightly because price increases are passed through more quickly than decreases.

²⁹For details on the construction of all three variables see Lewis (2009).

³⁰While we present results based on pass-through by the fifth day following a cost change, the regression results using pass-through rates on other days are very similar to those reported here.

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$	$ \Delta c = 2.69$	$ \Delta c = 4.92$
Price Cycle Present	.393** (.026)	.395** (.023)	.384** (.024)	.387** (.022)
Population per Station	.008 (.023)	.007 (.020)	-.015 (.027)	-.015 (.025)
Marketshare of Independent Stations	.133 (.080)	.128+ (.072)	.088 (.103)	.088 (.091)
HHI within Independent Stations	.249* (.099)	.231** (.089)	-.038 (.123)	-.038 (.110)
Population per Station \times (Price Cycle Present)			.002 (.051)	.003 (.025)
Marketshare of Independent Stations \times (Price Cycle Present)			.252 (.172)	.228 (.153)
HHI within Independent Stations \times (Price Cycle Present)			.789** (.206)	.736** (.183)
Cost Increase	.560** (.059)	.594** (.019)	.541** (.019)	.576** (.017)
Cost Decrease	.245** (.059)	.209** (.019)	.227** (.019)	.192** (.017)
Adjusted R ²	.960	.969	.960	.972
Number of Observations	142	142	142	142

Notes: “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. **, *, + indicate statistical significance at the 1%, 5%, and 10% levels respectively.

but it predicts much less of the overall variation in the speed of pass-through. Estimates in Columns 1 & 2 suggest that a one standard deviation increase in HHI within independents is, on average, associated with a 3 percentage point increase in the share of cost change passed through by the fifth day. Interestingly, Columns 3 & 4 reveal that the HHI within independents only helps to predict speed of response in cycling markets, not in non-cycling markets. Within markets with price cycles the estimates imply that a one standard deviation increase in the HHI within independents is associated with a 8 percentage point increase in the speed of pass-through.

Overall, the results reveal that the presence of price cycles is an important predictor of the speed of price response even after controlling for other observed market factors known to be

associated with the presence of cycles. Moreover, the presence of cycles has a much larger impact on the predicted response speed than any of the underlying market structure characteristics. While previous studies conclude that market concentration and factors associated with the ease of consumer search may be responsible for generating price cycles, our results show that these market factors do not *directly* cause large differences in the speed of price pass-through. It is only once a market has *tipped* into a cyclical price equilibrium, as a result of these or other factors, that it experiences a large increase in the speed of pass-through.

VI. 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 pass-through 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 studies 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 pass-through 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* 85:3 (2003), 772–776.
- Borenstein, Severin, Colin Cameron, and Richard Gilbert, "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?" *The Quarterly Journal of Economics* 112:1 (1997), 306–339.
- Cosslett, Steven and Lung-Fei Lee, "Serial Correlation in Latent Variable Models," *Journal of Econometrics* 27:1 (1985), 79–97.
- Deltas, George, "Retail Gasoline Price Dynamics and Local Market Power," *Journal of Industrial Economics* 56:3 (2008), 613–628.
- Doyle, Joseph, Erich Muehlegger, and Krislert Samphantharak, "Edgeworth Cycles Revisited," *Energy Economics* 32:3 (2010), 651–660.
- Eckert, Andrew, "Retail Price Cycles and the Presence of Small Firms," *International Journal of Industrial Organization* 21:2 (2003), 151–170.
- Engle, Robert F. and Clive W. J. Granger, "Co-Integration and Error Correction: Representation, Estimation, and Testing," *Econometrica* 55:2 (1987), 251–276.
- Lewis, Matthew, "Temporary Wholesale Gasoline Price Spikes have Long-lasting Retail Effects: The Aftermath of Hurricane Rita," *The Journal of Law & Economics* 52:3 (2009).
- "Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market," *Journal of Economics and Management Strategy* (Forthcoming).
- Maskin, Eric and Jean Tirole, "A Theory of Dynamic Oligopoly II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles," *Econometrica* 56:3 (1988), 571–599.
- Noel, Michael, "Edgeworth Price Cycles, Cost Based Pricing and Sticky Pricing in Retail Gasoline Markets," *Review of Economics and Statistics* 89:2 (2007a), 324–334.
- "Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market," *Journal of Industrial Economics* 55:1 (2007b), 69–92.
- "Edgeworth Cycles and Focal Prices: Computational Dynamic Markov Equilibria," *Journal of Economics and Management Strategy* 17:2 (2008), 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* 56:3 (2008), 581–612.
- Wang, Zhongmin, "Collusive Communication and Pricing Coordination in a Retail Gasoline Market," *Review of Industrial Organization* 32:1 (2008), 35–52.

Appendix

TABLE A1: NON-CYCLICAL ERROR CORRECTION MODEL COEFFICIENT ESTIMATES

	Positive Change		Negative Change	
$\Delta\text{Wholesale}_t$	0.131**	(0.007)	0.038**	(0.005)
$\Delta\text{Wholesale}_{t-1}$	0.083**	(0.006)	0.033**	(0.005)
$\Delta\text{Wholesale}_{t-2}$	0.048**	(0.006)	0.027**	(0.005)
$\Delta\text{Wholesale}_{t-3}$	0.062**	(0.006)	0.023**	(0.005)
$\Delta\text{Wholesale}_{t-4}$	0.048**	(0.005)	0.045**	(0.005)
$\Delta\text{Wholesale}_{t-5}$	0.056**	(0.006)	0.043**	(0.005)
$\Delta\text{Wholesale}_{t-6}$	0.045**	(0.005)	0.050**	(0.005)
$\Delta\text{Wholesale}_{t-7}$	0.014**	(0.005)	0.064**	(0.005)
$\Delta\text{Wholesale}_{t-8}$	0.000	(0.005)	0.059**	(0.005)
$\Delta\text{Wholesale}_{t-9}$	0.017**	(0.005)	0.023**	(0.005)
$\Delta\text{Wholesale}_{t-10}$	0.028**	(0.005)	0.011**	(0.005)
$\Delta\text{Wholesale}_{t-11}$	0.024**	(0.005)	0.025**	(0.005)
$\Delta\text{Wholesale}_{t-11}$	0.025**	(0.005)	0.031**	(0.005)
$\Delta\text{Wholesale}_{t-12}$	0.029**	(0.005)	0.033**	(0.005)
$\Delta\text{Wholesale}_{t-13}$	0.001	(0.005)	0.046**	(0.005)
$\Delta\text{Wholesale}_{t-14}$	-0.003	(0.005)	0.048**	(0.005)
$\Delta\text{Wholesale}_{t-15}$	0.004	(0.006)	0.033**	(0.005)
$\Delta\text{Wholesale}_{t-16}$	0.005	(0.005)	0.018**	(0.005)
$\Delta\text{Wholesale}_{t-17}$	-0.019**	(0.005)	0.021**	(0.005)
$\Delta\text{Wholesale}_{t-18}$	0.004	(0.005)	0.021**	(0.005)
$\Delta\text{Wholesale}_{t-19}$	0.010*	(0.005)	0.017**	(0.004)
$\Delta\text{Wholesale}_{t-20}$	0.032**	(0.005)	-0.006	(0.005)
⋮	⋮		⋮	
$\Delta\text{Retail}_{t-1}$	0.207**	(0.014)	0.195**	(0.026)
$\Delta\text{Retail}_{t-2}$	-0.067**	(0.013)	0.082**	(0.019)
$\Delta\text{Retail}_{t-3}$	0.003	(0.013)	-0.001	(0.019)
$\Delta\text{Retail}_{t-4}$	-0.012	(0.013)	0.020	(0.018)
$\Delta\text{Retail}_{t-5}$	-0.039**	(0.014)	-0.017	(0.020)
$\Delta\text{Retail}_{t-6}$	-0.038**	(0.012)	0.054**	(0.021)
$\Delta\text{Retail}_{t-7}$	0.054**	(0.012)	0.092**	(0.020)
$\Delta\text{Retail}_{t-8}$	0.008	(0.012)	-0.017	(0.019)
$\Delta\text{Retail}_{t-9}$	0.004	(0.012)	-0.075**	(0.018)
$\Delta\text{Retail}_{t-10}$	-0.037**	(0.011)	-0.001	(0.017)
⋮	⋮		⋮	
Error Correction Term	-0.006**	(0.003)	-0.070**	(0.003)
R^2			0.414	
obs			20313	

Notes: The dependant variable is the daily change in the retail price (ΔRetail_t). Robust standard errors are in parenthesis. **, *, + indicate statistical significance at the 1%, 5%, and 10% levels respectively.

TABLE A2: CYCLING MODEL COEFFICIENTS FROM THE PRICE AND SWITCHING EQUATIONS.

	Price Change Equation		Regime Switching Equation	
	Relenting	Undercutting	R → R	U → R
Retail _{t-1}	-.216** (.066)	-.0502** (.0020)	-.001 (.001)	-.249** (.016)
Wholesale _t	.176** (.071)	.0470** (.0020)	.007 (.002)	.254** (.016)
ΔWholesale _t	.104* (.052)	-.0103** (.0032)	-.055+ (.034)	.179** (.019)
ΔWholesale _{t-1}	.059* (.028)	.0074** (.0029)	.146** (.023)	-.106** (.018)
ΔWholesale _{t-2}	.114** (.045)	.0008 (.0027)	-.165** (.024)	-.056** (.014)
ΔWholesale _{t-3}	.195** (.042)	-.0033+ (.0022)	.032 (.025)	-.029* (.013)
ΔWholesale _{t-4}	.057 (.047)	.0017 (.0024)	.034+ (.026)	.101** (.014)
ΔWholesale _{t-5}	.092** (.032)	-.0015 (.0021)	.007 (.033)	.101** (.014)
ΔWholesale _{t-6}	-.031 (.018)	.0044* (.0024)	.012 (.017)	.025* (.015)
ΔWholesale _{t-7}	-.090** (.021)	.0063** (.0024)	-.028* (.016)	-.004 (.013)
ΔWholesale _{t-8}	.003 (.024)	.0120** (.0024)	-.076** (.019)	.037** (.011)
ΔWholesale _{t-9}	-.049 (.019)	.0119** (.0026)	.014 (.016)	-.062** (.013)
ΔWholesale _{t-10}	.110** (.044)	.0075** (.0026)	-.118** (.027)	-.053** (.012)
ΔWholesale _{t-11}	.070+ (.044)	.0044* (.0022)	.013 (.029)	.099** (.012)
ΔWholesale _{t-12}	.003 (.025)	-.0050* (.0024)	.119** (.022)	-.013 (.014)
ΔWholesale _{t-13}	.028+ (.019)	.0025 (.0023)	-.069** (.017)	.006 (.013)
ΔWholesale _{t-14}	.040* (.017)	-.0024 (.0026)	-.014 (.015)	.009 (.011)
ΔWholesale _{t-15}	.026 (.024)	.0041+ (.0026)	-.005 (.013)	-.060** (.013)
ΔWholesale _{t-16}	.130** (.033)	.0144** (.0026)	.060** (.022)	-.032** (.012)
ΔWholesale _{t-17}	.129** (.036)	.0030+ (.0021)	.005 (.022)	-.067** (.013)
Constant	10.5** (.564)	.1131 (.0801)	-.545+ (.357)	-.478* (.279)

Notes: **, *, and + indicate statistical significance at the 1%, 5%, and 10% levels.