UC San Diego
Recent Work

Title

Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market

Permalink

https://escholarship.org/uc/item/64j579g9

Author

Noel, Michael

Publication Date

2004

Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market

Michael Noel[*]

University of California, San Diego

December 18, 2003

Abstract

In this article, I exploit a new station-level, twelve-hourly price dataset to examine the strong retail price cycles in the Toronto gasoline market. The cycles are visually similar to the theoretical Edgeworth Cycles of Maskin & Tirole [1988]: strongly asymmetric, tall, rapid, and highly synchronous across stations. I test a series of predictions made by the theory about how firm behaviors would differentially evolve over the path of a cycle. The evidence is consistent with the existence of Edgeworth Cycles and inconsistent with competing hypotheses. One finding is that smaller firms are more likely than larger firms to initiate rounds of price undercutting but the reverse is true for rounds of price increases. While the cycles are an interesting phenomenon for study in their own right, the evidence has important implications for understanding market power in both cycling and non-cycling gasoline markets.

JEL Classification L13, L41, L81

1 Introduction

In many Canadian cities, retail gasoline prices are very volatile and in a unique and surprising way. Publicly available weekly price series in these cities exhibit a high-frequency cyclical pattern, asymmetric in nature, that begins with a large and sudden increase in retail prices that is generally complete in a week. Prices then fall gradually by small amounts over the following many weeks until the price is sufficiently near the wholesale price again that the prices jump back up and the cycle begins anew. This repeated pattern of behavior is strikingly
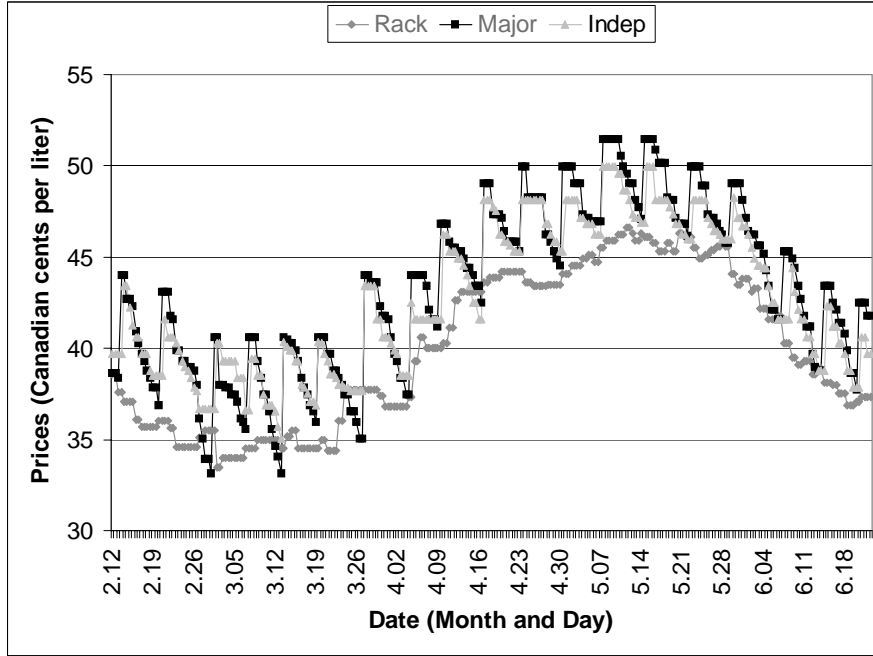
1

Figure 1: Retail Prices (Major, Independent) and Rack Price

similar in appearance to the theoretical (but in practice, arguably implausible) "Edgeworth Cycles" of Maskin & Tirole[1988].

In recent work, Noel[2003a] and Eckert[2003] examine a panel of Canadian retail gasoline markets to study this phenomenon. The authors use different estimation techniques to show that cycling activity is more prevalent in markets with more small firms, consistent with the theory of Edgeworth Cycles. Noel[2003a] goes on to show that relationships between the penetration of small firms and the structural characteristics of the cycle such as period, amplitude, and asymmetry are also consistent with the theory. Noel[2003a] argues that, contrary to common public perception, cycles are suggestive of increased competition.

Because these studies must rely on weekly, market-level data, they cannot observe individual pricing decisions of small and large firms along the cycle. This would be useful to know since the theory of Edgeworth Cycles make predictions about specific and differential firm behavior. With high-frequency, station-specific microdata, one could test if these interdependent behaviors are consistent with the theory of Edgeworth Cycles or with several other competing hypotheses.

While the cycles are an interesting phenomenon for study in their own right, evidence that they operate as Edgeworth Cycles would also have important

2

implications for understanding the degree of market power among the large multinational firms in cycling and non-cycling gasoline markets alike.

Therefore, I collect a new panel set of twelve-hourly retail gasoline prices for 22 service stations in the city of Toronto over 131 days in 2001.[1] I chose Toronto in part because of strong cycles suspected there. To my knowledge, no dataset with this fine a level of detail has been collected on cycles such as these before.[2]

The new dataset reveals a clean and asymmetric price cycle. Figure 1 shows the twelve-hourly price series for a representative station operated by a major integrated firm and one operated by an independent firm over the sample.[3] While the asymmetric pattern is clear in retail prices, there is none in the wholesale ("rack") price.

My first objective is simply to demonstrate the existence of a high-frequency cycling phenomenon in this market econometrically and parametrically describe the structural characteristics of the cycles. I show that a Markov switching regression model, adapted from Cosslett & Lee[1985] and Ellison[1994], is well suited to this general problem. I find that the cycles are tall relative to markups, with an amplitude of 5.6 cents per liter on average. This represents 170% of the average retail markup and 364% of the average markup at the bottom of a cycle. Amplitudes of 10 cpl (equivalent to 24.5 U.S. cents per gallon) – an increase that generally occurs in a single instant – were common. I also find that the period of the cycle is about a week but is not anchored to any particular day of the week.

The most interesting finding is the asymmetry of price movements. Each station tends to increase its price by the full height of the cycle in a single jump and yet lowers its price in small amounts over the following four to ten days. At first glance, stations also appear to act in close synchronicity.

My second objective is to show that not only do these cycles appear like Edgeworth Cycles but that they *act* like Edgeworth Cycles as well. In particular, I isolate the pricing behaviors of small independents and large integrated firms and compare these behaviors to those predicted by the theory of Edgeworth Cycles (Maskin & Tirole[1988] and Eckert[2003]). I find that each type of firm is important in generating cycles and in different ways. Consistent with the theory, I find that undercutting tends to be triggered by smaller firms while new rounds of marketwide price increases are more likely triggered by larger firms. Period-by-period pricing behaviors are also consistent. The results are inconsistent with other hypotheses for the existence of asymmetric cycles, such as shifting demand, rack discounting, or changing station inventories.

I follow with a short discussion of competition in cycling markets vis a vis non-cycling markets. Even without the benefit of variation in market structure, I can argue that a penetration of small independent firms into a market, alongside

---

[1] I thank my wife Cheryl Noel for collecting the data with a tape recorder while driving a roundabout route to and from work each day.

[2] Websites that periodically post prices reported by consumers on a daily or weekly basis are notoriously unreliable and do not occur on a station specific basis at regular intervals.

[3] Taxes have been removed. There are excise taxes of 24.7 cpl and a sales tax of 7%.
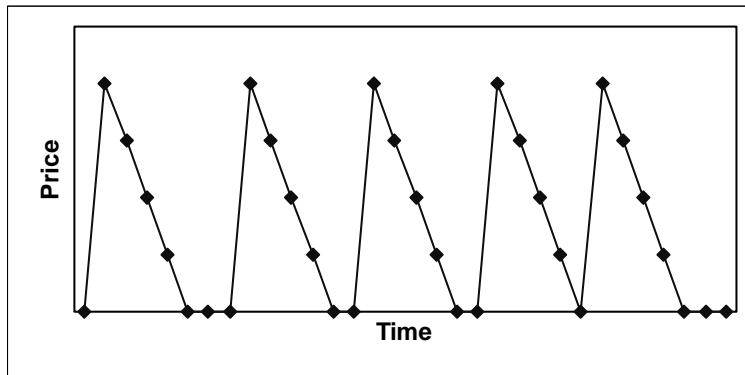
Figure 2: Edgeworth Cycle

the few large ones that exist (in all markets), are associated with an increased likelihood of cycling behavior. Cycles are suggestive of more active competition – non-cyclical pricing suggests relatively weaker competition. While the general volatility and synchronicity of prices in cycling markets has led to many claims of price fixing, the asymmetric nature of the cycle is not as obvious to the public and has not been well understood.

Section 2 discusses the related literature and Section 3 lays out my empirical framework. A short discussion of the data is in Section 4. In Section 5, I report results on cycle characteristics and in Section 6 I report my results on the influence of large and small firms. Section 7 discusses alternative hypotheses and policy under a price cycle, and Section 8 concludes.

2 Theory and Literature

The price cycles observed in these cities are very similar in appearance to the theoretical "Edgeworth Cycles" of Maskin & Tirole[1988]. In their paper, Maskin & Tirole consider a dynamic Bertrand price setting game under which two identical firms produce homogenous goods under constant costs. Firms set prices alternately along a finite price grid, and each responds to its opponent's action from the previous period. Firms split demand if prices are equal and the lower priced firm captures the entire market if not equal. The authors show that even under identical supply and demand conditions there are two distinct sets of Markov perfect Nash equilibria (MPE). The first is the familiar set of focal price equilibria, and the second, the authors call "Edgeworth Cycles."

In the longer downward portion of an Edgeworth Cycle, firms repeatedly undercut one another by one notch on the grid in order to steal market share. When price reaches marginal cost, each firm has a positive probability of raising its price back to the "top" of the cycle. A war of attrition results as each waits

for the other to go first. When one firm resets the price, the other immediately follows and price undercutting begins anew. Figure 2 shows an example of a time path of cycling prices in their model. The strong cyclical pattern observed in Toronto appears consistent with this type of dynamic equilibrium.

One limitation of the Maskin & Tirole analysis is that the two firms are identically sized. Eckert[2003] extends the theoretical model to the case of two firms of different sizes, where a firm of a larger "size" simply means it receives a larger fraction of consumers when prices are identical. When prices are unequal, it is assumed the lower priced firm (whether it is the large or small firm) can and does serve the entire market itself. As a result, the smaller the firm, the greater the incentive to undercut. The author concludes that when the two firms are sufficiently different in size, only Edgeworth cycling can exist. Relevant to this paper, the author also shows that along the cycle the small firm will undercut on each of its moves but the large firm will instead price match on some or all of its moves. That is, the small firm leads the large firm down the cycle. Conversely, for reasonable parameter values, the large firm is more likely than the small firm to be first to increase price back to the top of the cycle.

In Noel[2003b], in which I extend the Maskin-Tirole and Eckert frameworks along numerous dimensions including fluctuating marginal costs, degree of differentiation, and number of firms, I show that coordination problems can occur in cycle resetting when there are more than two firms. This further increases the likelihood that large firms (who control the price for many stations) will emerge as reliable cycle resetters.

The second objective of this paper will be to show that individual firms behave in these ways, consistent with the theory of Edgeworth Cycles, but inconsistent with other explanations.

The setting of the models of Edgeworth Cycles lends well to the Toronto retail gasoline market. Gasoline is relatively homogeneous, frequently purchased, and firm level demand is highly-elastic.[4] Retail prices of regular gasoline are displayed on tall billboards – easy for consumers to compare and easy for competing firms to monitor. Menu costs are absent: in Toronto, retail prices can change several times a day. Discussions with regional managers also suggest that an alternating moves game – where each firm monitors and then responds to changes by other firms – is an appropriate description of behavior.[5]

Retail gasoline competition has been of longstanding interest to economists,

---

[4] High firm-level price elasticity is an important factor. Imperial Oil reports claim that many consumers do respond to differences as low as 0.2 cents per liter. (Majors generally only price in odd decimals so 0.2 is the minimum undercut.) Perhaps additional utility is being gained from paying the lowest price, since savings would only be about a dime on a fillup. My own anectodal evidence when collecting this data suggests that a difference of 1 cpl at two nearby stations (*very* rare and very brief) indeed has a large impact on volumes.

[5] However, gasoline retailers are capacity constrained. A single small firm with few retail outlets can steal only a negligible fraction of the market by undercutting and is easily ignored by the large firm. Only when there are many many such firms, individually small but collectively large, will more widespread undercutting steal enough market share to induce a response from the large firm. In Toronto, approximately half of the stations are operated by small independents.

policymakers, and consumers and empirical studies for markets in which cycles have not been detected are many. Approaches have included comparing estimated markups to oligopoly models using a specific price war episode (Slade[1987,1992])[6], examining responses to wholesale price shocks (Borenstein, Cameron & Gilbert[1997]), mergers (Hastings & Gilbert[2001]), and the business cycle (Borenstein & Shepard[1996]), and examining multiproduct station pricing (Shepard[1991]). Few papers however have specifically addressed asymmetric price cycles of this nature.

For the United States, Allvine & Peterson[1974] note similar patterns in some western U.S. cities in several episodes in the 1960s and early 1970s. They argue that the sudden halt to price cycling in these cities in 1972 may have been the result of suffocation of wholesale supplies to independent retailers in advance of the oil shortage of 1973, but no empirical analysis is performed. In a short note, Castanias & Johnson[1993] note the Edgeworth-like appearance of the cycles in Los Angeles from 1968 to 1972, and present simple summary statistics on price changes and runs.

For Canada, Eckert[2002] uses weekly average price data for the city of Windsor from in the early 1990s, to show that the nature of the price cycle results in rack price increases being passed through to retail prices more quickly than decreases. This contrasts earlier studies by Hendricks[1996], Lerner[1996], and Godby et. al[2000] that included non-cycling cities. While the large literature on asymmetric rack-retail passthrough in the U.S. and elsewhere (Borenstein, Gilbert, & Cameron[1997] and many others) typically assume reversion to a single long-run steady-state retail price and are conducted for markets that are known *not* to exhibit retail price cycles, the existence of cycles suggests a new potential source for passthrough asymmetry.

Eckert[2003] motivates his theoretical model (described above and found in that paper) with some interesting correlations between overall price rigidity and concentration ratios in Canadian markets. The rigidity variable is defined as the fracton of times weekly average spot prices did not change from week to week within each year within a given city, and the concentration ratios are based on year-end station counts for 19 cities and 6 years.

Noel[2003a] explicitly models three distinct pricing patterns in these markets – cycles, sticky pricing, and cost-based pricing – in a 19 city, 11 year sample and builds estimates the prevalence of each pattern and the structural characteristics of the cycle using estimated parameters from a Markov switching regression. Price data is weekly and station counts are bimonthly. The author finds that with a greater penetration of small firms, price cycles (as different from volatile prices) are more prevalent, sticky prices less prevalent, and the cycles have

---

[6]Slade[1987, 1992] examines a price war in Vancouver in the summer of 1983 but it is different than the repeated, high-frequency price cycles analyzed here. The price war appears isolated in an era of stable prices and is postulated to have occurred because of an unanticipated demand shock. In those papers, the author reports evidence of tacit collusion and found that rivals respond more quickly to price increases by a major firm than to decreases while the opposite was true for reactions to independents. The literature review here to papers focuses on repeated, asymmetric Edgeworth-like price cycles.

shorter periods, greater amplitudes, and are *less* asymmetric: prices rise as quickly but fall much faster. Each of these relationships is consistent with the theories of Edgeworth Cycles. A welfare analysis shows that markups are also lower during cycling periods than during non-cycling periods.

Two limitations of the data used in these papers is its weekly frequency and that it is not station-specific. Cycles may be partially masked from the econometrician if averaging across stations obscures sudden price increases or if cycles have periods less than a few weeks. The dataset I use in this paper is free of these issues and allows a uniquely clean look at individual firm behaviors. Moreover, I focus on the city of Toronto specifically because of the strong cycles expected to exist there.[7]

3   Empirical Framework

The first objective of this paper is to show how to build parametric measures of the characteristics of the cycle for use later on. I do this by taking nonlinear transformations of the parameters produced from a Markov switching regression.

I allow the regimes to be station-specific – that is, each station can potentially follow a cycle of its own. Since the data suggests that only cycling activity is occuring, I model a station as operating under one of two regimes at a given point in time. They are:

1. the relenting phase (regime "R"),

2. the undercutting phase (regime "U"),

and I later model the switching probabilities between the two.

A nice feature of the Markov switching regression is that it allows serial correlation in the estimated regimes. The true underlying regimes are considered unobservable since price movements in different regimes can in principle look identical. For example, a zero price change or small price increase (decrease) by a station may still be considered a part of its undercutting phase (relenting phase) depending which regime we believe the station was in the previous period and the estimated switching probabilities. A regular switching regression, which relies only on information contained in the current observation, cannot do this.

Of course, with this new dataset, the cycle is clear enough that one could get some similar results by separately analyzing price increases and price decreases or by using a regular switching regression. If measuring characteristics were the only concern (which is not the case here), one could even get away with just eyeballing the data or examining the summary statistics. However, the Markov switching regression framework is preferable for several reasons. First,

---

[7]Both from the weekly data used in Noel[2003a] and my own experience. The Toronto price series shows up in weekly data as very volatile week-to-week price changes. While the period is bottomcoded at two weeks and the amplitudes are underestimated, the model in Noel[2003a] is able to categorize cycles in 84% of periods overall, and in more than 98% of periods in each of five of the last seven years.

it is general and the techniques described here continue to apply to data that is not as clean.[8] Second, it is less ad hoc: no assumptions need be made about how to categorize, for example, zero price changes or small price increases in the middle of extended periods of price decreases. Imposing minimum or maximum cutoffs for inclusion into a particular regime would otherwise produce estimates influenced by subjective categorization. Third, since it directly estimates the probability of switching between regimes, I can derive intuitive formulae for the characteristics of the cycle and easily allow those characteristics to covary with variables of interest, all within a single specification.

3.1 The Regimes

The portion of the cycle during which I anticipate finding prices that rise sharply in a short time I call the relenting phase and the portion during which I anticipate finding prices that gradually fall I call the undercutting phase. Consider a station s at time t which is operating under regime i. I assume that the firm who operates station s sets its retail price according to the function

$$\Delta RETAIL_{st} = \begin{cases} X_{st}^i \beta^i + \varepsilon_{st}^i & \text{with prob. } 1 - \gamma_{st}^i \\ 0 & \text{with prob. } \gamma_{st}^i \end{cases} \tag{1}$$

where $\Delta RETAIL_{st} = RETAIL_{st} - RETAIL_{s,t-1}$ and $RETAIL_{st}$ is the retail price, $\left(X_{st}^i\right)'$ is an $K^i \times 1$ vector of explanatory variables, β^i is a $K^i \times 1$ vector of parameters and ε_{st}^i is a normally distributed error term with mean zero and variance σ_i^2. Let $\alpha^i = E(\Delta RETAIL_{st} \mid X_{st}^i)$.[9]

Especially important with very high frequency data that produces many zero price changes, I have allowed in each regime a mass point at zero. Also note that the regime specifications are built identically and no restrictions are placed on the sign of the price change for inclusion in a given regime.

Setting the X^i to a vector of ones, for example, allows a simple estimation of the average price changes in each phase of the cycle and contributes to measures of its vertical characteristics. When I include firm type dummies and cycle position variables into X^i in section 6, I can study differential pricing behaviors along the path of the cycle.[10]

3.2 The Switching Probabilities

There are four Markov switching probabilities in total (from each of two regimes to each of two regimes.) Let I_{st} be equal to "R" and "U" when station s at time t is in the relenting phase regime and undercutting phase regime respectively.

---

[8] For example, in Noel[2003a].

[9] Rather than first price differences on the LHS, one can model the relenting phase using a price level on the left hand side and the rack price on the right hand side all with similar results. The results will show that relenting firms tend to reset price to achieve a standard markup rather than a fixed price level.

[10] Particulars of each specification are discussed together with results in later sections.
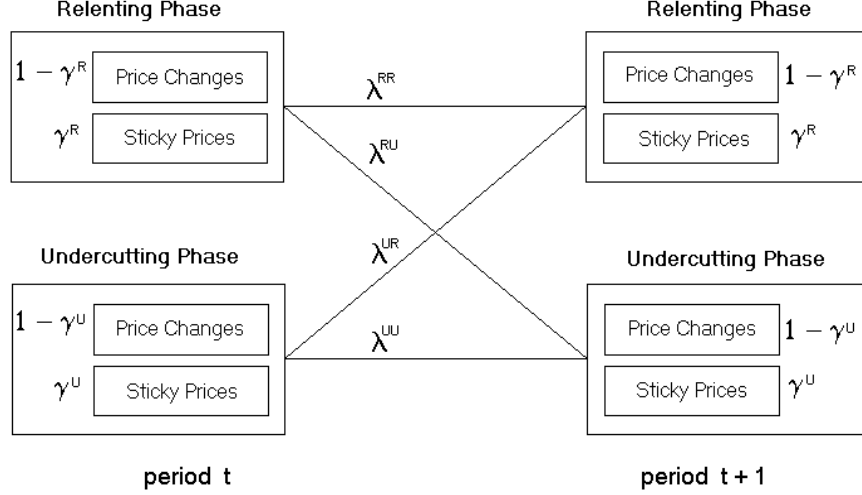
Figure 3: Regimes and Switching Probabilities

Then the probability that a station switches from regime i in period $t-1$ to regime "R" in period t is given by the logit form:

$$
\begin{aligned}
\lambda_{st}^{iR} &= \Pr(I_{st} = \text{``}R\text{''} \mid I_{s,t-1} = i, W_{st}^i) \\
&= \frac{\exp(W_{st}^i \theta^i)}{1 + \exp(W_{st}^i \theta^i)}, \quad i = R, U
\end{aligned}
$$

and $\lambda_{st}^{iU} = 1 - \lambda_{st}^{iR}$, $i = R, U$ to satisfy the adding up constraint. Let Λ_{st} be the 2×2 switching probability matrix whose ij^{th} element is λ_{st}^{ij}. Each $(W_{st}^i)'$ is an $L^i \times 1$ vector of explanatory variables that affects the switching probabilities out of regime i and θ^i is an $L^i \times 1$ vector of parameters.

In addition, let J_{st}^i be the indicator function equal to 1 when, conditional on operating under regime i, the price at that station does not change. Then the probability that the station's price will not change in any given period, conditional on regime i, is modeled as the logit probability:

$$
\Pr(J_{st}^i = 1 \mid I_{st} = i, V_{st}^i) = \gamma_{st}^i = \frac{\exp(V_{st}^i \zeta^i)}{1 + \exp(V_{st}^i \zeta^i)} \tag{2}
$$

where $(V_{mt}^i)'$ is a $Q^i \times 1$ vector of explanatory variables and ζ^i is an $Q^i \times 1$ vector of parameters.

Figure 3 outlines the structure of the model.

Setting the W^i to a vector of ones, for example, yields estimates for average switching probabilities which I use to measure the horizontal characteristics of the cycle. When I include firm type dummies and cycle position variables into W^i in section 6, I can study differential regime-switching behavior along the

9

path of the cycle. Similarly, including firm type dummies and cycle position variables into V^i allows for differential "holding" patterns along the path of the cycle.

The core model parameters $(\beta^i, \theta^i, \zeta^i)$ in each specification are simultaneously estimated by the method of maximum likelihood. Numerical methods are used to calculate robust Newey-West standard errors on the core estimates. Estimates of the switching probabilities, cycle features, and all partial derivatives are derived by joint non-linear transformations of the core parameter estimates and standard errors by the multivariate delta method.

3.3 The Anatomy of a Price Cycle

By combining the switching probabilities and the within-regime parameters, I can now derive intuitive formulae for the structural characteristics of the cycles directly from the parameters of the model. There are three dimensions of interest – cycle period, amplitude, and asymmetry.

The expected period of the cycle, or the distance between consecutive troughs, is just the expected length of a relenting phase plus the expected length of an undercutting phase. Since the probability of "switching" from regime i in period $t-1$ to the same regime again in period t is λ^{ii}, I derive the expected duration of regime i as

$$E(duration\ of\ regime\ i) = \frac{1}{1 - \lambda^{ii}} \tag{3}$$

and the expected period of the cycle as

$$E(period) = \frac{1}{1 - \lambda^{RR}} + \frac{1}{1 - \lambda^{UU}} \tag{4}$$

To derive the amplitude of the price cycle, I multiply the expected duration of the relenting phase with the expected relenting phase price change. One could also the undercutting phase to calculate the vertical fall (rather than the vertical rise) and the long term stationarity of prices over the sample period ensures these measures are about the same. Therefore, I calculate expected amplitude as

$$E(amplitude) = \frac{(1 - \gamma^R)\alpha^R}{1 - \lambda^{RR}} \ or \ \frac{-(1 - \gamma^U)\alpha^U}{1 - \lambda^{UU}} \tag{5}$$

One of the most interesting characterisics of the cycles is their asymmetry. There are two dimensions on which to measure this: horizontally and vertically. I define "horizontal asymmetry" as the ratio of the duration of the undercutting phase to the duration of the relenting phase:

$$E(horizontal\ asymmetry) = \frac{1 - \lambda^{RR}}{1 - \lambda^{UU}} \tag{6}$$

and "vertical asymmetry" as the (negative of the) ratio of the average price change in an relenting phase to the average price change of the undercutting

phase:

$$E(vertical\ asymmetry) = \frac{-(1-\gamma^R)\alpha^R}{(1-\gamma^U)\alpha^U} \tag{7}$$

Again, the long run stationarity of prices ensures this measures to be roughly the same.

4 Data

I collect and use a unique dataset of twice-daily retail prices for 22 service stations along major city routes in east Toronto over 131 days between February 12^{th} and June 22^{nd} 2001. The stations I surveyed are a representative mix of large major national and regional firms and smaller independent firms. Thirteen of the stations surveyed are operated by major national or regional firms, nine by independents.[11] Twelve firms are represented in total including all major national and regional firms.

Retail prices, $RETAIL_{st}$, are for regular unleaded, 87 octane, self-serve gasoline, in Canadian cents per liter. The descriptive specifications of Section 5 use after-tax prices (since firms compete on these); the behavioral specifications of Section 6 use tax-exclusive prices (relevant for profit margins.) Taxes are almost entirely lump sum and results are unaffected by this choice. The wholesale price I use is the daily spot rack price for the largest wholesaler at the Toronto rack point, $RACK_{st}$, as collected and reported by MacMinn Petroleum Advisory Service.[12] Ancillary data such as firm and station characteristics, source of price control and inventory deliveries were self-collected.

Summary statistics for rack and retail prices are shown in table 1. The US\$/gallon price equivalents are US\$1.08/gallon before tax, US\$1.78/gallon after tax, and an average rack-retail markup of US\$0.08/gallon.

Before discussing results, I mention a few minor data issues. First, when firms raise price, they do so in such close synchronicity that even with twice-daily data I will identify a leading group of stations rather than a single leading station or firm. Several continuous surveillances were carried out and reaction times were found to be as fast as a few hours by some followers.

Another issue involves the structure of the gasoline industry and the choice of wholesale prices. In the gasoline industry, the largest firms are integrated into wholesaling and retailing – and such a firm will operate some of its retail outlets itself while it leases others to private dealers. In urban centers, the former is dominant and for all major firm branded stations in my sample, the head office controls prices. For these stations, no wholesale market transaction takes place.

---

[11] Majors are defined as those that are integrated into refining and retailing, independents are retailers only. A major firm will generally have a much larger retail presence than an independent firm.

[12] A single wholesale price was used to ensure averages did not mask large jumps in the wholesale price. There is no substantive difference between using a firm-specific rack prices or daily spot averages.

There are also many independent retailers, individually small with no refining division, and it is these firms that buy at the rack price.

Overall, the rack price is the best available measure of the wholesale price. In contrast to the United States, there is little discounting off posted rack prices in Canada, and any small discounts that exist are not tied to movements in the retail price.[13] Lerner[1996] among others further argues the rack price measures the wholesaler's opportunity cost of wholesale gasoline sold to dealers. Because of readily available U.S. sources of wholesale gasoline constraining local wholesalers (for example, see Hendricks[1996]), the rack price is reasonably modeled as exogenous to retail price setting.

5   Description of the Toronto Retail Price Cycle

I now return to my first objective and begin with a simple empirical description of the anatomy of the Toronto price cycle. To do this, I first need to introduce the basic specifications and present the within-regime parameter estimates and switching probabilites from which to build that description.

5.1   The Regimes

In specification (1), the expected price changes in each of the two cycle regimes ($\alpha^i = E(\Delta RETAIL_{st} \mid X^i_{st})$, $i = R, U$), all four switching probabilities (λ^{ij}, discussed below) and the probability of sticky pricing conditional on being in regime i, (γ^i), are assumed constants. This simple setup allows a single measure of each descriptive feature of the retail price cycle.

In table 2, I report estimates from each within-regime equation. In specification (1), the expected retail price change in the undercutting phase is -0.75 cpl (conditional on a non-zero change) and I find sticky prices in 43% of undercutting periods. Since the data is twelve-hourly, this shows a station undercuts a little more than once per day.

In contrast, the expected retail price change in a relenting phase is 5.57 cpl and sticky prices are effectively non-existent.

In specification (2), I separate stations into two types – those operated by large major (national or regional) firms and those operated by small independents. I allow the α^i, γ^i, and λ^{ij} to depend on station type. The table shows that a relenting independent station increases its price by a significantly smaller amount than does a major, 5.19 cpl versus 5.78 cpl. The average undercut is only slightly smaller in absolute value for independents but independents tend to hold prices sticky more often than a major.

---

[13] The reader will note that the markup, as measured by the current price less the posted rack price can go slightly negative on occasion.

5.2   Switching Probabilities

The switching probabilities for each specification are reported at the bottom of table 2.

In specification (1), I find that a station that has relented in the period before will continue in a relenting phase one period (12 hours) later with less than a 1% probability ($\lambda^{RR} = 0.008$). With 99% probability, it begins a new undercutting phase.

A station currently undercutting will continue to undercut in the next period with a 92% probability ($\lambda^{UU} = 0.92$) and has only an 8% chance of switching to a price-restoring relenting phase ($\lambda^{UR} = 0.08$).

In specification (2), where majors and independents are estimated separately, I find that the estimates are effectively the same across station type. In fact, the same results hold even when each firm or each station is estimated separately.

5.3   Anatomy of a Price Cycle

Plugging in the regime parameter estimates and switching probabilities into the formulae above, I now summarize the structural characteristics of the cycles and report them in table 3.

In specification (1), I estimate the expected period of the cycles at 13.78 half-day periods, or a little less than a week. This consists of an average relenting phase lasting 1.01 periods (half a day) and an undercutting phase lasting 12.78 periods (6.39 days.)[14] That the period is about a week suggests that perhaps cycles are driven by day-of-the-week supply or demand changes. I will consider and ultimately dispel the calendar effects hypothesis and other alternative explanations later on.

This result also shows that the relenting phase is quite short. My anectodal observation is that, for a given station, it is instant. The full height of the cycle is achieved in a single price increase.

The expected amplitude of the cycles is 5.61 cpl and 5.48 cpl under the relenting and undercutting phase measures respectively.[15] Since the average ex-tax retail price in the sample is 43.05 cpl and the average markup is 3.31 cpl, the expected amplitude (and one-time relenting phase increase) is 13% of the average ex-tax price, 170% of the average markup overall, and 364% of the average retail markup just prior to a relenting phase. One-time price increases of 10 cents per liter (equivalent to 24.5 US cents per gallon) were common throughout the sample.

Finally, I report an estimate of horizontal asymmetry as 12.68 – greater than one at a very high level of significance. The undercutting phase is over twelve times longer than the relenting phase. The estimate of vertical asymmetry is 13.00, again significantly greater than one.

---

[14] Half a day is the frequency of the data: the relent at a given station is generally instant.

[15] As the two measures of amplitude are never significantly different from one another in any given specification, I report only the relenting phase calculations in the table.

Differentiating between stations operated by majors and by independents in specification (2), I find virtually identical cycle periods and asymmetries within each type. This is not surprising if the two sets of cycles are interrelated, but this alone does not show synchronicity of the cycles. I also find significant differences in amplitudes across station type: 5.83 cpl for majors and 5.23 cpl for independents.

These descriptive results confirm the existence of cycles that are rapid, tall, and asymmetric in the direction we would expect if they were in fact Edgeworth Cycles. They point to a potentially strong interrelationship between the two, but only hint at potential differences. In the next section, I move beyond the simple description and to my second objective and main focus of the paper.

6   Small Firms vs. Large Firms

Edgeworth Cycles are an interesting theoretical construct, but can this be the phenomenon we are observing in practice? Could other alternative hypotheses generate similarly asymmetric and high-frequency cycles instead? To answer this question, I turn to an examination of the behaviors of small and large firms at different points along the path of cycle. The theory of Edgeworth Cycles makes specific predictions about how the groups interact and how their pricing behaviors differentially evolve along the path of the cycle. My second objective is to test if the observed firm behaviors are consistent with Edgeworth Cycles and inconsistent with alternative hypotheses.

I allow for changing behavior along the path of the cycle using two key variables in specification (4): $POSITION$ and $FOLLOW$, described below. To allow differential behavior by firm size, I then interact all RHS variables including $POSITION$ and $FOLLOW$ with large (major) firm and small (independent) firm dummies.

Define $POSITION$ as the difference between the lagged retail price and the current rack price, less taxes, $RETAIL_{s,t-1} - RACK_{st} - TAX_{st}$. This is intended as a measure of the position of a station's ex-tax price relative to the bottom of its cycle (approximated by marginal cost.) Since I want to test for changes in the aggressiveness of a firm's pricing strategy based on its stations' position within the cycle (and differentially by firm size), I allow the expected price change in each regime (α^i) and the probability of sticky pricing in the undercutting phase (γ^U) to vary with $POSITION$.[16] This is done by including $POSITION$ in the X^R, X^U, and V^U matrices.

Changes in $POSITION$ should also influence regime change. As a given station nears the bottom of its cycle, one expects an increasing probability that a firm will switch a station out of its undercutting phase and into its relenting phase. Thus I include $POSITION$ in the switching probability out of the undercutting phase $(\lambda^{UR}$ via $W^U)$, but not out of the relenting phase since two consecutive periods of relenting are extremely rare in the data $(\lambda^{RR} \sim 0)$. For examples of switching probabilities out of the undercutting phase at various levels

---

[16]Previous specifications show sticky prices are effectively non-existent in relenting phases.

of $POSITION$, see table 4.[17] The probability of switching from undercutting to relenting ramps up quickly as $POSITION$ falls.

The dummy variable $FOLLOW$ is intended to capture differential behavior of large and small firms in the transition from undercutting to relenting. I am interested both in how large and small firms self-select into roles as leaders and followers in cycle resetting and also how their behaviors differ conditional on their roles. Let $FOLLOW_{st}$ be equal to one in period t if some other station has already relented as of the previous period but station s still has not. Since all stations relent each time and relenting rounds are well separated, this variable (and more complex versions of it) is easily constructed.[18] Once all have relented, $FOLLOW$ is set back to zero for every station. Since I will want to test for differences in the aggressiveness of pricing strategies at the very bottom of the cycle, I allow the probability of switching from undercutting to relenting (λ^{UR}) and the expected price change in the subsequent relenting phase (α^R) to depend on $FOLLOW$.

6.1   Undercutting

Assume that I have already shown that the cycles for all the stations are highly synchronous and that all stations have set just completed their relenting phases. Prices are relatively high and stations are together at the tops of their respective cycles. From this point, the theory of Edgeworth Cycles suggests that smaller firms have a greater incentive to initiate a new round of price undercutting than larger firms. If we observe more active undercutting by *smaller* firms near the tops of the cycles, it would be consistent with the theory. Other hypotheses described later, such as calendar or inventory effects, do not predict differential behavior by firm size.

The main results of the paper are found in table 5. In the top half, I report partial derivatives of the expected price changes (α^i) and the probability of sticky prices during undercutting phases (γ^U) with respect to $POSITION$, and of the expected relenting phase price change with respect to $FOLLOW$.[19] Each is interacted with firm size and reported separately for small (independents) and large (major) firms. I discuss the switching probabilities in the bottom half below.

---

[17]This is based on a specification identical to (2) but that includes $POSITION$ in X^R, X^U, V^U, and W^U. The example is for major firms although that for independents is similar. (Once the $FOLLOW$ variable is added, there would be four such matrices, two for majors and two for independents.) The partial derivative is calculated by $\frac{\partial \lambda^{UR}}{\partial POSITION} = \theta^U_{POSITION} * \lambda^{UR} * \lambda^{UU}$.

[18]I also estimated the model using the number or fraction of firms who have previously relented (simple or weighted according to distance or discounted over time) and results are similar. It would be computationally infeasible to estimate a fully specified state model where each station is in one of 2^{22} states (depending on who has and has not yet relented). It is quite unlikely that stations are concerned with the full distribution of which individual stations have and have not relented at a point in time, however.

[19]The partial derivatives are calculated as $\frac{\partial \alpha^i}{\partial y} = \beta^i_y$ and $\frac{\partial \gamma^U}{\partial POSITION} = \zeta^U_{POSITION} * \gamma^U * (1 - \gamma^U)$.

15

My results are consistent with the theory of Edgeworth Cycles. As predicted, I find that small independents are more aggressive near the top of their cycles. There is little difference between the two types in terms of the how the size of their undercuts change from top to bottom but there is in terms of how their "holding" patterns change.

Near the top of the cycles and early in the undercutting phase, actual price undercutting (of any size) is substantially more prevalent among the smaller firms while sticky prices are more prevalent with large firms. Nearer to the bottom, the roles reverse and we see that undercuts are more common with larger firms and sticky prices more common with smaller firms ($\frac{\partial \gamma^U_{MAJ}}{\partial POSITION} = 0.036$, $\frac{\partial \gamma^U_{IND}}{\partial POSITION} = -0.034$). Each estimate is significantly different from zero as well as from each other.

This shows that small independents are more likely to initiate new undercutting phases. It also suggests that large firms try to support higher top-of-the-cycle prices for awhile but ultimately chase small firms downwards as the price gap between them grows too large. In fact, just prior to a new round of relenting, the prices at majors are often *below* those of the independents. This differential behavior of large and small firms is not consistent with alternative hypotheses.

How the size of the undercuts themselves change along the cycle is also consistent the theory. While invariant to firm size, actual undercuts (as opposed to sticky prices) become very slightly smaller ($\frac{\partial \alpha^U_i}{\partial POSITION} = -0.04$) as firms move further down the cycle. While the straight theory has that each firm undercuts by an identically-sized minimum amount until marginal cost is reached and the war of attrition begins, the empirical implication is that firms become less aggressive in undercutting nearer the cycle bottom. This result suggests that firms when they undercut do so less aggressively, albeit slightly, close to the bottom.

6.2 Relenting

The theory of Edgeworth Cycles suggests that larger firms have a greater incentive and greater coordinating ability to trigger a new round of relenting phases once markups become too low. Moreover, reactions are so fast that the cycles across stations should be closely synchronous. If we observe earlier relenting activity by *larger* firms near the bottom of the cycle that is followed very quickly with relenting activity by smaller firms, it would be consistent with the theory. While some alternative explanations might predict close synchronicity, they do not predict a close timing pattern by firm size.

In the bottom half of table 5, I calculate and report estimated switching probabilities (λ^{ij}) by firm size and by $FOLLOW$ status for several relevant values of $POSITION$. This presentation will be more intuitive to the reader than reporting the partials that underlie them. I consider $POSITION = 1.53$, the average value prior to a relent, and $POSITION = 0$, commonly reached in the data. At time t, a "follower" is simply a station whose $FOLLOW$ dummy

16

flag is on – that is, it has not yet relented but at least one other station has and a marketwide return to higher prices is underway. A "leader" is a station whose $FOLLOW$ dummy is off – no station had just relented and each is still a *potential* leader in terms of cycle resetting.

The results are again consistent with the theory of Edgeworth Cycles. I find that larger firms are more likely than smaller firms to initiate new rounds of relenting phases. This contrasts the earlier finding that smaller firms are more likely than larger firms to initiate new rounds of undercutting phases. I also find that reactions of following firms are very fast, and the larger the firm the faster is the reaction.

Consider the switching probabilities at $POSITION = 1.53$, the average value just prior to a new relenting phase, and consider first the $LEADER$ columns. Conditional on no station having yet relented, the probability that a large major firm will switch a station into its relenting phase in the current period is 9.1%. The corresponding value for a small independent is only 2.6%. The estimates are statistically different from each other at much better than the 1% level of significance. This evidence shows that large firms are much more likely to be in the leading group of stations than small firms.[20]

Next consider the $FOLLOWER$ columns. Conditional on at least one station having relented in a previous period, the probability that a large major will switch a station into a relenting phase in the current period is 93%. The corresponding value for a small independent is 72%. These estimates translate into two more important results.

First, the probabilities are high. Firms large and small respond extremely quickly to a previous relenting phase by switching into relenting phases themselves, most likely within the following half-day. This produces highly synchronous cycles across stations.

Second, the estimate for majors is statistically and significantly greater than that for independents at much better than the 1% level. Since all follower stations eventually relent during each round, the estimates show that majors respond even more quickly on average than do independents. Independents occasionally delay more than one twelve-hour period (28% probability), but it is rare for a major to do so (7%).

The transition from undercutting phases to relenting phases is now more clear. Large firms are more likely initiate new rounds of relenting, any remaining large firms are more likely to follow first and finally small firms tend to follow behind. Reaction times are very fast overall and cycles appear highly synchronous.[21] My findings are consistent with the theory of Edgeworth Cycles and are inconsistent with alternate hypotheses I consider later.

I also report the case when $POSITION$ reaches 0 cents still with no station relenting. The probability that a major branded station relents as part of the leading group in the current half-day rises to 23.5%. For an independent, it is

---

[20] Recall that because of fast reaction times, I typically identify a leading group of stations rather than a single leading station even with twelve-hourly data.

[21] With price data that was just tri-daily or even bi-daily, stations would appear to be in perfect synchronicity.

17

6.6%. Once some station relents, competing stations follow in the very next period with a probability close to one, majors (97%) even more likely than independents (87%). The same relationships hold for any value of $POSITION$.

The last two minor results are also consistent with Edgeworth Cycles. In the theory, when a following firm relents, it will set a price just below that of the leading firm, effectively making the first undercut. I find evidence of this effect. Near the top of the table, I report that a major firm who follows tends to raise price by 0.2 cents less than if it had led. The corresponding effect for independents is insignificant, likely due to the fact that independents individually are probably all followers but the fastest ones can make it into the leading group.

Finally, in an Edgeworth Cycle, a firm will reset its price near the monopoly price, which in turn a function of marginal cost. I find that a station's price change during a relenting phase will be greater the closer it had been to the bottom of the cycle. The coefficient on α^R for majors and independents are fairly close to one ($\frac{\partial \alpha^R_{MAJ}}{\partial POSITION} = -0.91, \frac{\partial \alpha^R_{IND}}{\partial POSITION} = -0.87$), suggesting a roughly standard markup is being reinstated with each relent.

6.3   A Fuller Specification of Price Changes

The previous parsimonious specification (4) captures the main results of the paper. In the final specification (5), I add to that specification by modeling the price responses of large and small firms, leaders and followers, in a little more detail. I show how the current markup and expected future costs influence the magnitude of price changes. I also show that the cycles of stations are interdependent, with firms setting prices in response to prices of other firms. The results of this section are again consistent with the theory of Edgeworth Cycles and inconsistent with other explanations.

First, consider the relenting phase of the cycle. I earlier showed that the lower the position of a station just prior to relenting, the higher tends to be the price increase, reinstating a roughly standard markup. Moreover, if rack prices are expected to rise in the near future, one would expect to see relenting phase price increases to be relatively greater in anticipation.

Optimally, I would like to show that changes in cycle position and expected future costs affect the magnitude of price increases by leaders but not those by followers. Followers, I postulate, simply react to leaders' price increases by slightly undercutting the new higher prices, irrespective of cycle position or expected costs. The econometric problem with respect to followers is that the strong synchronicity in the cycles (fast reaction times relative to position or cost changes) means these two variables are highly collinear with the $MAXDIFF$ variable discussed below. Therefore, I will impose the restriction that cycle position and expected future costs do not influence follower behavior directly and I focus instead on how they affect leader behavior.

Let $\Delta_8 RACK_{st}$ be the difference between the current rack price and the

expected rack price four days (eight periods) into the future.[22] I include the interactions $POSITION_{st}*(1-FOLLOW_{st})$ and $\Delta_8 RACK_{st}*(1-FOLLOW_{st})$ in the relenting phase price change equation (in X^R).

As mentioned, I expect relenting phase price changes by followers to closely follow price changes by leaders. The greater the difference between a station's last period price and the maximum price in the market, the greater should be its relenting phase price change.[23] Define $MAXDIFF_{st} = RETAIL_{s,t-1} - MAX(RETAIL_{r,t-1})$ where the maximum is taken over all stations $r \neq s$. I include $MAXDIFF_{st}*FOLLOW_{st}$ in the relenting phase price change equation (in X^R) to model follower behavior. This measure is not meaningful for leaders (since the maximum price will be at a station that is still undercutting relatively low prices).

Leader and follower main effects were also included in the relenting phase price change equation and all variables were interacted with the firm size dummy to allow differential behavior between large and small firms.

Now consider the undercutting phase. If truly undercutting, a station should lower its price more often and by larger amounts (in absolute value) when the differential between its previous price and the market minimum is greater. Define $MINDIFF_{st} = RETAIL_{s,t-1} - MIN(RETAIL_{r,t-1})$, $r \neq s$.[24] I include $MINDIFF_{st}$ and $POSITION_{st}$ into the expected price change equation (in X^U) and the probability of sticky pricing in the undercutting phase (in V^U), and interact them with the firm size dummy.

I report my results in table 6. As before, I find major firms leading a new round of relenting phases increase prices by an additional 0.918 cpl for each cpl the pre-relenting phase markup is lower ($\frac{\partial \alpha^R_{MAJ}}{\partial(POSITION*(1-FOLLOW))} = -0.918$.). I also find that leading majors do raise prices by more in anticipation of rack price increases ($\frac{\partial \alpha^R_{MAJ}}{\partial(\Delta_8 RACK*(1-FOLLOW))} = 0.162$.). Estimates for independents are similar but less precise due to the small number of independents in the leading group.

For followers in the relenting phase, I show they raise prices by more when the absolute price difference between their prices and the market maximum is greater. ($\frac{\partial \alpha^R_{MAJ}}{\partial(MAXDIFF*FOLLOW)} = -0.950$, $\frac{\partial \alpha^R_{IND}}{\partial(MAXDIFF*FOLLOW)} = -1.176$.) Recall that the price change equations are linear and include a constant, so

---

[22] For future rack prices, I use a simple forecasting procedure based on current and lagged rack prices and parameters estimated from the full sample. The particular choice of model has little effect. Results are similar if I use forecast based on a model estimated with only prior data. Results are also similar using other lead lengths or the difference between current and lagged racks directly.

[23] More specifically, a follower would care about the minimum of the prices of those stations that have already relented. Because the reactions times are so fast, either measure produces very similar results. I choose the $MAXDIFF$ measure since it does not need assumptions on how to classify a station in the rare case that its relenting phase lasts two periods.

[24] I also estimate a version in which the min function is taken over only those competing stations on the same "side" of the latest round of relenting phases as the given station. This would mean stations undercutting from the top of the cycle are concerned only with the prices of other stations that have relented as well. Results are similar but slightly stronger in this case.

while the effect of $MAXDIFF$ on following independents is relatively greater, the expected price change for independents remains much lower than that for majors $(\alpha_{MAJ}^R > \alpha_{IND}^R)$ and the new price set by independents is always below the maximum.

Turning to the undercutting phase, I find that firms do react to changes in each others' prices along the path of the cycle, consistent with the theory. Both large and small firms will lower its station prices more often and by greater amounts in absolute value when its price is relatively higher above the market minimum. Notice that the results for non-zero price changes are similar across firm size, but the result for the probability of sticky pricing (γ^U) is significantly lower in absolute value for majors. This stems from the earlier finding that majors try to hold prices constant at the top even as the minimum market price falls, offsetting an otherwise stronger negative effect (as seen for independents). That earlier finding, by the way, continues to hold strong: the effect of cycle position on sticky prices is significantly positive for majors and significantly negative for independents.

In conclusion, the interplay between large major firms and small independents is consistent with the theory of Edgeworth Cycles. I find that larger firms are more likely than smaller firms to initiate new rounds of relenting phases, but the opposite is true for undercutting phases. The magnitude of relenting phase price increases are sensitive to changes in cycle position and expected future costs. Reactions of following firms are very fast, and the larger the firm the faster is the reaction. The pricing behaviors of firms in the undercutting phase are also closely interrelated, resulting in cycles that are highly synchronous.

7 Discussion

7.1 Alternative Hypotheses

The fineness of the data has allowed me to test for numerous relationships and behavioral differences by firm size as predicted by the theory of Edgeworth Cycles. It also enables me to reject other leading alternative hypotheses.

For example, one possible explanation for the cycling phenomenon is shifting demand. That the cycles are highly synchronous with a duration of roughly a week suggests that its shape may be the result of predictable day-of-the-week demand changes. However, this hypothesis is quickly dispelled. First, relenting phases do not occur on the same day each week but rather are common all throughout the week.[25] On any specific calendar day, prices may be high or low week to week.

Second, this argument requires that demand increase sharply one day (a different day each week) and then fall by small increments each following day, which seems highly unlikely given personal experience. While leading firms may

---

[25] The exception is on weekends, when industry activity is low. The probability of relenting on each day of the week is: Monday 18%, Tuesday 32%, Wednesday 22%, Thursday 18%, Friday 10%, Saturday and Sunday 0%.

prefer a marketwide relent earlier in the week to take advantage of high weekday sales, it can only lead if margins are sufficiently low to induce others to follow. Most weeks see one round of relenting, although weeks with two and weeks with none are not uncommon in this market. I also note that in national data (see Noel[2003a]), cycle durations vary from weeks to months depending on the city.

Third, there is little reason in a demand story to suspect differential timing of relenting phases or undercutting intensity by firm size. A station of a given size should equally likely be a leader or follower in a round or relenting. All stations should be equally likely to begin undercutting. In order to match the data, demand at major firms would have to spike up about a day before independents, stay at its peak a day or so after demand for independents started to fall, and then fall relatively more rapidly than independents after that. A demand story also cannot explain the influence of future rack prices on amplitudes or why following majors reset prices just below that of leading majors.

On a related note, I find that there is no "long weekend effect" – that is, firms do not set prices higher specifically for the long weekend. This is in contrast to a recent government study which claims to have found one and cites it as evidence of non-competitive behavior.[26] Part of the problem is that while price jumps occur roughly once a week, the increase in the week prior to each long weekend is routinely singled out in the media. In reality, those jumps are not exceptional, and can occur on any day of the previous week.[27]

A second alternative hypothesis is that station refuelling and inventory depletion would generate a price cycle. This is unlikely on many fronts. Given limited underground tank capacity and uncertain demand, and for now assuming an exogenously determined delivery schedule, firms would set price just low enough to exhaust inventory prior to the next delivery.[28] To generate a cycle, firms would have to repeatedly overestimate or repeatedly underestimate demand as the cycle progresses. In the former, firms slowly adjust price downward over time and in the latter, they adjust upward. Alternatively, if the firm is risk averse to running out of inventory (and cannot hold extra), it may start with higher than expected-market-clearing prices and lower them gradually as refuelling day draws near and sales uncertainty falls.

First, it is difficult to argue that firms will systematically over- or underestimate demand all the time. Having said that, this would produce cycles whose asymmetries would go different ways for different firms, which we do not observe. Demand uncertainty so great and risk aversion of firms so strong that prices change by as much as 25% over the course of a week is also implausible, especially when delivery schedules are endogenously set and can absorb any risk of depletion.

Second, since stations replenish inventories on different days, it cannot explain the synchronicity of relenting phase price increases across all stations in

---

[26]Ironically, a upward shift in demand in a competitive market would have the same qualitative effect.

[27]Government of Canada[1998].

[28]This also requires that deliveries are infrequent enough so that firms would run out at competitive prices.

the market. Obviously, then, it cannot also explain the specific timing of different sized firms in relenting phases, and cannot explain the specific timing of firms and the interrelationship of prices on the way down.

Another possibility is that discounts off the posted rack price, unobserved to the econometrician, create a rack price cycle that accounts for the retail price cycles. Again, this explanation does not work. Rack price discounts are much smaller (generally less than a cpl) than the amplitude of the cycle and vary by station size but not over time as required for a cycle. (It also seems strange that a firm would symmetrically change its own posted rack prices up and down only to offset it by adjusting the discount.) Again, this hypotheses cannot explain large and small firm timing on the cycle.

I conclude that my set of findings are inconsistent with these alternative hypotheses of cycle generation. They are consistent with the theory of Edgeworth Cycles.

7.2   Welfare and the Price Cycle

With this new data, the retail price cycle is clear: the price at a given station increases by a large amount suddenly and then falls a little each day. While consumers are accustomed to prices that change daily, it is interesting that few are aware that virtually all the small changes are price decreases. In an informal poll I conducted of 58 people living in neighborhoods near the sample stations in June 2001, the average respondent believed that a station's price would change about once a day. Conditional on changing, it was believed on average 58% were price *increases*. In the actual sample, prices do change about once a day but of those 86% are price *decreases*. While the larger one-time price jumps are widely reported in the press and shape consumers' perceptions, the many small offsetting price decreases that follow are less noticed. A popular impression is that gasoline prices are "always going up", in reality they are almost always going down. Even over the sample period where prices actually fell overall, there were renewed calls for industry regulation.

The existence of a retail price cycle argues that there is a substantial, although not perfect, degree of competition currently in the Toronto retail sector. If firms were perfectly colluding, we might see stable prices near what is now the top of the cycle. Competitive forces can keep pulling the prices back down. Noel[2003a] shows that in other gasoline markets where few small firms exist, prices are stable and margins are higher. And to the extent that consumers time their purchases in cycling markets (my anectodal observation is that few do), they are better off still.[29] Of course, where price cycling exists, some market power is required for a firm to effectively reset the cycle.

That the existence of a price cycle is indicative of competition in the marketplace is in contrast to classic views on price wars. A traditional and symmetric price war is often seen as a punishment mechanism to facilitate collusion. In

---

[29]Perhaps more elastic consumers, with a lower opportunity cost of time, would benefit most.

these repeated, asymmetric gasoline price wars, the price cuts appear not to be punishments nor to hold any signalling value.

8   Conclusion

In this paper, I have presented evidence that the retail price cycles are a dominant feature of the Toronto market, and that there is differential behavior among small firms and large firms along it. My findings are consistent with the theory of Edgeworth Cycles and inconsistent with other hypotheses.

Using a Markov switching regression framework, I found cycles have an expected amplitude of 5.6 cents and an expected period of about a week. The cycles are strongly asymmetrical: prices increase suddenly and then fall about three-quarters of a cent each day over the next six or seven days.

I find that larger firms are more likely than smaller firms to initiate new rounds of relenting phases, but the opposite is true for undercutting phases. The magnitude of relenting phase price increases are sensitive to changes in cycle position and expected future costs. Reactions of following firms are very fast, and the larger the firm the faster is the reaction. The pricing behaviors of firms in the undercutting phase are also closely interrelated, resulting in cycles that are highly synchronous.

While studying the mechanisms that generate the cycles is an interesting in its own right, my results give insight into the competitive intensity in these markets. While some market power is required to reset prices higher in cycling markets, pro-competitive forces in the form of small undercutting firms helps bring them back down. These small-firm forces appear absent in non-cycling markets.

Finally, I am not yet aware of similar Edgeworth-like cycles currently in other product markets. However, the interested observer can look across markets for certain characteristics that may lend well to cycling activity. For example, markets with many small firms alongside at least one large firm capable of coordinating marketwide price increases are more likely to experience cycles. The product itself would likely be relatively homogeneous, highly elastic at the firm level, non-durable (in practice if not in principle), and frequently purchased. Consumer switching costs and firm menu costs should also be close to zero. Whether evidence is found for the existence of Edgeworth Cycles in other product markets remains to be seen.

References

[1] Allvine, F. & J. Patterson. "Highway Robbery, An Analysis of the Gasoline Crisis", Bloomington: Indiana University Press, 1974, pp. 192-197, p.243.

[2] Borenstein, S. "Selling Costs and Switching Costs: Explaining Retail Gasoline Margins", RAND Journal of Economics 22, 1991.

[3] Borenstein, S., A. Shepard. "Dynamic Pricing in Retail gasoline Markets", RAND Journal of Economics 27, 1996.

[4] Borenstein, S., A. C. Cameron, R. Gilbert. "Do Gasoline Markets Respond Asymmetrically to Crude Oil Price Changes?" Quarterly Journal of Economics 112, 1997.

[5] Castanias, R. and H. Johnson. "Gas Wars: Retail Gasoline Price Fluctuations", Review of Economics and Statistics 75, 1993.

[6] Cosslett, S. and L. Lee. "Serial Correlation in Latent Variable Models", Journal of Econometrics 27, 1985.

[7] Eckert, A. "Retail Price Cycles and Response Asymmetry", Canadian Journal of Economics 35, Feb. 2002.

[8] Eckert, A. "Retail Price Cycles and Presence of Small Firms", International Journal of Industrial Organization 21, 2003.

[9] Ellison, G. "Theories of Cartel Stability and the Joint Executive Committee," RAND Journal of Economics 25, Spring 1994.

[10] Government of Canada. "Report of the Liberal Committee on Gasoline Pricing in Canada." Ottawa, June 1998.

[11] Hendricks, K. Analysis and Opinion on Retail Gas Inquiry. Report prepared for Industry Canada, 1996.

[12] Hastings, J. and R. Gilbert. "Market Power, Vertical Integration, and the Wholesale Price of Gasoline", mimeo, 2003.

[13] Imperial Oil Limited. http://www.esso.ca/news/issues/mn_price.

[14] Lee, L. and R. Porter. "Switching Regression Models with Imperfect Sample Separation Information – With an Application on Cartel Stability," Econometrica 52, 1984.

[15] Lerner, G. "Evaluation of the Six Resident's Allegation of Price-Fixing in the Canadian Petroleum Industry," Report prepared for the Competition Bureau, Industry Canada, 1996.

[16] MJ Ervin & Associates. "Canadian Retail Petroleum Markets Study." Report prepared for Industry Canada, 1997.

[17] Maskin E., and J. Tirole. "A Theory of Dynamic Oligopoly II: Price Competition, Kinked Demand Curves and Edgeworth Cycles," Econometrica 56, 1988.

[18] Noel, M. "Edgeworth Price Cycles, Cost-based Pricing and Sticky Pricing in Retail Gasoline Retail Markets," UCSD mimeo, 2003a.

[19] Noel, M. "Edgeworth Price Cycles and Focal Prices: Computational Dynamic Markov Equilibria" UCSD mimeo, 2003b.

[20] Porter, R. "Optimal Cartel Trigger Price Strategies," Journal of Economic Theory 29, 1983.

[21] Shepard, A. "Price Discrimination and Retail Configuration," Journal of Political Economy 99, 1991.

[22] Shepard, A. "Contractual Form, Retail Price, and Asset Characteristics in Gasoline Retailing," RAND Journal of Economics 24, 1993.

[23] Slade, M. "Interfirm Rivalry in a Repeated Game: An Empirical Test of Tacit Collusion," Journal of Industrial Economics 35, 1987.

[24] Slade, M. "Price Wars in Price-Setting Supergames," Economica 56, 1989.

[25] Slade, M. "Vancouver's Gasoline Price Wars: An Empirical Exercise in Uncovering Supergame Strategies," Review of Economic Studies 59, 1992.

[26] Slade, M. "Strategic Motives for Vertical Separation: Evidence from Retail Gasoline Markets," Journal of Law, Economics, and Organization 14, 1998.

[27] United States Energy Information Administration, "U.S. Downstream Independents Acquire National Prominence in the 1990s," http://www.eia.doe.gov/emeu/finance/sptopics/restructure/highlite4.html.

Table 1: Rack and Retail Price Changes and Runs

	Mean	Std.Dev.	Minimum	Maximum
RETAIL (price before tax)	43.09	4.44	32.2	51.8
RETAIL (price after tax)	72.53	4.75	60.9	81.9
RACK	39.77	4.00	33.5	46.0
POSITION	3.31	2.08	-3.2	7.6
MAXDIFF	-1.39	1.30	-10.6	2.4
MINDIFF	1.56	1.72	-5.1	10.6

In Canadian cents per liter.

Table 2: Within-Regime Results and Switching Probabilities

Specification	(1)	(2)	
	pooled	majors	indeps
RELENTING PHASE (dep. var. $= \Delta RETAIL_{st}$)			
$\alpha^R = E(\Delta RETAIL_{st} \mid I_{st} = \text{``R''})$	5.576	5.782	5.192
(average price change)	(0.083)	(0.101)	(0.142)
σ_R	1.650	1.610	1.655
	(0.066)	(0.085)	(0.100)
$\gamma^R = \Pr(J_{mt} = 1 \mid I_{st} = \text{``R''})$	0.000	0.000	0.000
(fraction sticky prices)	(0.000)	(0.000)	(0.000)
UNDERCUTTING PHASE (dep. var. $= \Delta RETAIL_{st}$)			
$\alpha^U = E(\Delta RETAIL_{st} \mid I_{st} = \text{``U''})$	-0.751	-0.767	-0.720
(average price change)	(0.008)	(0.010)	(0.015)
σ_U	0.459	0.467	0.441
	(0.009)	(0.012)	(0.012)
$\gamma^U = \Pr(J_{mt} = 1 \mid I_{st} = \text{``U''})$	0.429	0.422	0.441
(fraction sticky prices)	(0.007)	(0.008)	(0.014)
SWITCHING PROBS. ($\lambda^{ik} = \Pr(I_{st} = k \mid I_{s,t-1} = i)$)			
λ^{RR} ($R \rightarrow R$)	0.008	0.008	0.007
	(0.004)	(0.005)	(0.007)
λ^{RU} ($R \rightarrow U$)	0.992	0.992	0.993
	(0.004)	(0.005)	(0.007)
λ^{UR} ($U \rightarrow R$)	0.078	0.078	0.078
	(0.001)	(0.002)	(0.003)
λ^{UU} ($U \rightarrow U$)	0.921	0.921	0.921
	(0.001)	(0.002)	(0.003)

Standard errors in parentheses calculated by delta method.

Table 3: Key Cycle Characteristics

Specification	(1)	(2)	
	pooled	majors	indeps
HORIZONTAL FEATURES			
Relenting Phase Duration	1.008	1.008	1.007
	(0.004)	(0.007)	(0.007)
Undercutting Phase Duration	12.780	12.779	12.784
	(0.291)	(0.358)	(0.484)
Cycle period	13.788	13.787	13.792
	(0.291)	(0.358)	(0.485)
Horizontal Asymmetry	12.680	12.677	12.687
	(0.291)	(0.363)	(0.491)
VERTICAL FEATURES			
Cycle Amplitude	5.619	5.828	5.232
	(0.082)	(0.098)	(0.144)
Vertical Asymmetry	13.001	13.053	12.897
	(0.224)	(0.275)	(0.390)

Durations and frequency in terms of half-day periods, amplitude in cents
per liter, measures of asymmetry are unit free. Standard errors in
parentheses calculated by numerical methods.

Table 4: Behavioral Effects of Cycle Position

Specification:	(3)		
	majors		indeps
$\frac{\partial \lambda^{UR}}{\partial POSITION}$	-0.065		-0.051
(eval. at average pre-relent POSITION)	(0.006)		(0.006)

EXAMPLE SWITCHING MATRICES (for major firms)

POSITION =	6.68	3.31	1.53	0	-3.24
	(average	(average	(average		(minimum)
	post-relent)	overall)	pre-relent)		
λ^{UR}	0.003	0.046	0.110	0.272	0.803
	(0.001)	(0.002)	(0.006)	(0.022)	(0.038)
λ^{UU}	0.996	0.953	0.889	0.727	0.197
	(0.001)	(0.002)	(0.006)	(0.022)	(0.038)

Standard errors in parentheses calculated by delta method.

Table 5: Leaders and Followers

Specification:	(4)	
	majors	indeps
$\frac{\partial \alpha^R}{\partial FOLLOW}$	-0.200	0.084
	(0.104)	(0.096)
$\frac{\partial \alpha^R}{\partial POSITION}$	-0.897	-0.873
	(0.027)	(0.061)
$\frac{\partial \alpha^U}{\partial POSITION}$	-0.042	-0.034
	(0.007)	(0.027)
$\frac{\partial \gamma^U}{\partial POSITION}$	0.028	-0.034
(eval. at average overall POSITION)	(0.002)	(0.005)

SWITCHING PROBS. (Undercutting Phase)

$POSITION = 1.53$ (average pre-relent)

	LEADERS			FOLLOWERS	
	majors	indeps		majors	indeps
λ^{UR}	0.091	0.026		0.925	0.720
	(0.006)	(0.003)		(0.045)	(0.045)
λ^{UU}	0.919	0.974		0.075	0.280
	(0.006)	(0.003)		(0.045)	(0.045)

$POSITION = 0$

	LEADERS			FOLLOWERS	
	majors	indeps		majors	indeps
λ^{UR}	0.235	0.066		0.974	0.868
	(0.020)	(0.011)		(0.016)	(0.030)
λ^{UU}	0.765	0.934		0.026	0.132
	(0.020)	(0.011)		(0.004)	(0.030)

Standard errors in parentheses calculated by the delta method.

Table 6: Response to Other Stations' Prices.

Specification	(5)	
	majors	indeps
$\frac{\partial \alpha^R}{\partial(POSITION*(1-FOLLOW))}$	-0.918	-0.856
	(0.028)	(0.141)
$\frac{\partial \alpha^R}{\partial(\Delta_s RACK*(1-FOLLOW))}$	0.162	0.186
	(0.320)	(0.123)
$\frac{\partial \alpha^R}{\partial(MAXDIFF*FOLLOW)}$	-0.950	-1.176
	(0.042)	(0.053)
$\frac{\partial \alpha^U}{\partial POSITION}$	-0.003	-0.020
	(0.005)	(0.007)
$\frac{\partial \alpha^U}{\partial MINDIFF}$	-0.185	-0.174
	(0.018)	(0.021)
$\frac{\partial \gamma^U}{\partial POSITION}$ (eval. at average POSITION)	0.041	-0.026
	(0.004)	(0.006)
$\frac{\partial \gamma^U}{\partial MINDIFF}$ (eval. at average POSITION)	-0.016	-0.072
	(0.009)	(0.021)

Standard errors in parentheses calculated by the delta method.
Constants not reported. Switching probabilities similar to
previous table and are not repeated here.