EDGEWORTH PRICE CYCLES: EVIDENCE FROM THE TORONTO RETAIL GASOLINE MARKET*

MICHAEL D. NOEL†

I exploit a new station-level, twelve-hourly price dataset to examine the strong retail price cycles in the Toronto gasoline market. The cycles appear similar to theoretical Edgeworth Cycles: 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. While the cycles are an interesting phenomenon for study in their own right, the evidence has important policy and welfare implications.

I. INTRODUCTION

IN MANY CANADIAN CITIES, RETAIL GASOLINE PRICES follow a high-frequency, asymmetric price cycle. Publicly available weekly price series show the cycle begins with a large and sudden increase in retail prices followed by many small price decreases over subsequent periods. Once markups are sufficiently small, prices jump back up and the cycle begins anew. The repeated pattern of behavior is strikingly similar in appearance to the theoretical (but in practice, arguably implausible) 'Edgeworth Cycles' of Maskin & Tirole [1988]. As discussed below, Maskin & Tirole derive their Edgeworth Cycles as a Markov equilibrium outcome of a dynamic homogeneous-good Bertrand game where firms alternate in choosing prices. But is the cycling phenomenon observed in Canadian retail gasoline markets really Edgeworth Cycles?

Pricing dynamics in these markets are not well understood, yet understanding the mechanism driving the cycle is important in many ways. From a policy perspective, evidence suggestive of Edgeworth Cycles would help rule out other hypotheses such as covert collusion, the subject of numerous

resource-consuming investigations. It would teach us how welfare-enhancing reversion back to low market prices after every peak works and what causes prices to peak again. It would also teach us that prices, while volatile, are to some extent predictable. Thus better informed, elastic consumers could increase their surplus by appropriately timing purchases of gasoline, which could in turn generate Coasian dynamics and lower prices.

Evidence of Edgeworth Cycles would also be interesting in its own right, as such cycles have seldom been documented empirically. Their existence would contrast assumptions of reversion to a single, long run, steady state price made in most intertemporal models of gasoline markets. Combined with quantity data, one could also estimate short run own and cross price elasticities over a wide range of prices on the equilibrium path within a short controlled time frame. However, to date, very little empirical work has been done to understand these cycles.

Two relevant articles are Noel [2007] and Eckert [2003] who examine weekly market prices across Canadian retail gasoline markets, some of which cycle and others of which do not. However, these studies, discussed in the literature review, have two important limitations. First, the weekly spot-average nature of the data obscures the finer cycle details and makes it impossible to test a central structural prediction of the theory. The theory states that the full height of the cycle at a particular station should be achieved in a single price increase and then be followed by a consecutive sequence of small price decreases.

Second, the theory of Edgeworth Cycles makes specific predictions about large and small firm behavior and how those behaviors differentially evolve along the path of the cycle. The three main behavioral predictions made by the theory are: (1) firm reactions are very fast but not simultaneous, (2) small firms tend to lead prices downwards, and (3) large firms tend to lead prices upwards. The market level data used in these studies cannot test these predictions of the Edgeworth Cycle model.

With high-frequency, station-specific data, however, one can. In this article, I present a new dataset of twelve-hourly retail gasoline prices for 22 service stations in the city of Toronto over 131 consecutive days in 2001. I chose Toronto in part because conversations with industry insiders suggest cycles are fastest there. In Figure 1, I show 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.[1] While the asymmetric cycle is clear in retail prices, there is not one in the wholesale ('rack') price.

Using a Markov switching regression model, adapted from Cosslett & Lee [1985] and Ellison [1994], I briefly parameterize and estimate average or

---

[1] Taxes have been removed. There are excise taxes of 24.7 Canadian cents per liter (cpl) and a sales tax of 7%.
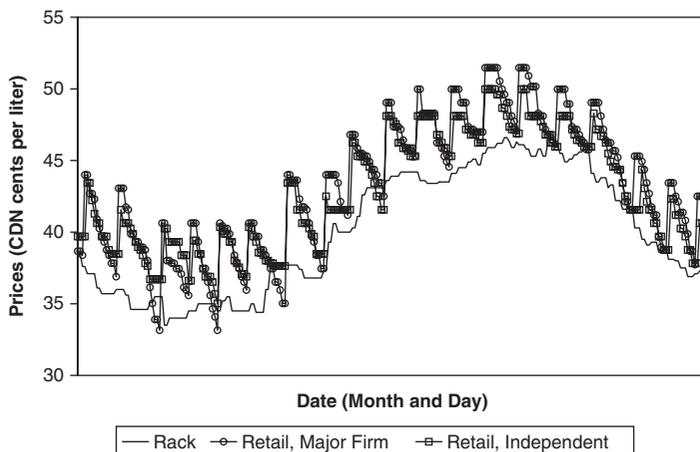
Figure 1
Retail Prices (Major Firm, Independent Firm) and Rack Price

'typical' cycle characteristics. In support of the structural prediction of the theory, I find that each station tends to increase its price by the full height of the cycle in a single jump but lowers its price in small amounts over four to ten days. At first glance, stations also appear to act in close synchronicity.

The main exercise of this article is then to identify the competitive process driving the cycles and show it is consistent with the theory of Edgeworth Cycles. I isolate the pricing behaviors of small independents and large integrated firms and test the three behavioral predictions outlined above. I find support for each of these predictions. The results are inconsistent with other hypotheses for the existence of asymmetric cycles, such as shifting demand, asymmetric discounts from the rack price, changing station gasoline inventories, or covert collusion.

In Section II, I discuss the theory and literature and in Section III, I present my empirical framework. A short discussion of the data is in Section IV. In Section V, I report estimates to describe a typical retail price cycle and in Section VI, I turn to a competitive analysis of large and small firms. Section VII contains a discussion of competing hypotheses and Section VIII concludes.

II.   THEORY AND LITERATURE

The price cycles observed in Toronto and in other Canadian cities are very similar in appearance to the theoretical 'Edgeworth Cycles' introduced by Edgeworth [1925] and formalized by Maskin & Tirole [1988].

Consider the following extension of the Maskin & Tirole [1988] model. Two infinitely-lived profit-maximizing firms compete in a dynamic pricing

game by alternately setting prices. Once set, the price for that firm is fixed for two periods. Prices are chosen from a discrete price grid. Marginal cost, c_t, is also allowed to vary over time, and is chosen by nature from a discrete cost grid under a uniform distribution for simplicity. Each firm earns current period profits of

$$(1) \qquad \pi_t^i(p_t^1, p_t^2, c_t) = D^i(p_t^1, p_t^2) * (p_t^i - c_t)$$

where D^i is the demand for firm i.

The strategies of each firm are allowed to depend only on the payoff-relevant state in each period, i.e., they are Markov. In this case, the state variables are the opponent's price from the previous period and current marginal cost. Let Firm 1's value function, in a period in which it is the active price setter but prior to learning current marginal cost, be

$$(2) \qquad V^1(p_{t-1}^2) = E_c\left(\max_{p_t}\left[\pi_t^1(p_t, p_{t-1}^2, c_t) + \delta_i W^1(p_t)\right]\right)$$

where

$$(3) \qquad W^1(p_{s-1}^1) = E_c\left(E_{p_s}\left[\pi_s^1(p_{s-1}^1, p_s, c_s) + \delta_i(V^1(p_s))\right]\right)$$

and similar for Firm 2. The discount factor for Firm i is δ_i. Each firm, when active, sets price to maximize the present discounted value of its future profit stream, or V^i without the outside expectation.

The Maskin & Tirole [1988] model can be recovered from this setup by setting $\delta_1 = \delta_2$, $c_t = c$ for all t and D^i is the standard homogenous Bertrand demand function. The authors show two different types of equilibria are possible: focal price equilibria and 'Edgeworth Cycle' equilibria. In an Edgeworth Cycle, firms repeatedly undercut one another to steal the market (the 'undercutting phase'), until price reaches marginal cost. At that point, a war of attrition ensues with each firm mixing between raising price and remaining at marginal cost.

While the length of the undercutting phase is not certain, the length of the 'relenting phase' is. The price increases at a given station in a single period, before undercutting starts again. The model thus predicts a clear asymmetric shape (the structural prediction) and extremely fast but not simultaneous reactions (behavioral prediction (1)). The model does not make a clear prediction about amplitudes, however. The top of the cycle price may be above or below the monopoly price and many amplitudes are possible in equilibrium. An example of an Edgeworth Cycle for two symmetric firms is shown in Figure 2.

Eckert [2003] extends the Maskin & Tirole analysis to allow for firms of different size. While maintaining the assumptions that $\delta_1 = \delta_2$ and $c_t = c$, the author models D^i as standard Bertrand except that firms split the market unequally at equal prices. The author shows that the smaller firm (with the
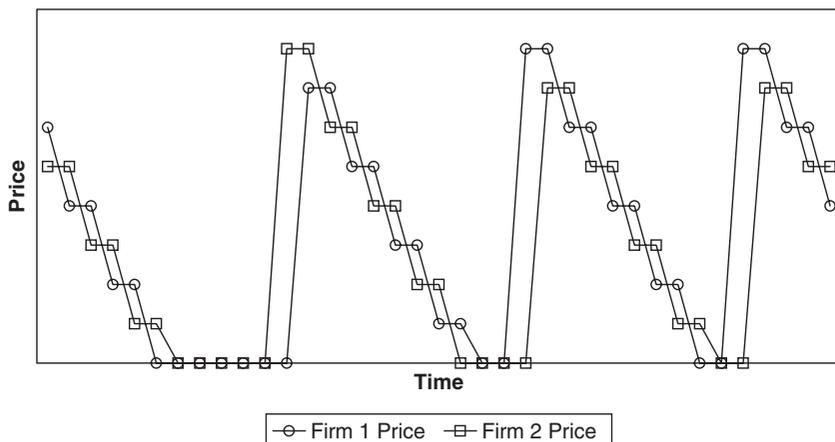
Figure 2
Theoretical Edgeworth Cycle
Firms are symmetric in this example with equal and constant marginal cost. The top of the cycle
is at a price that may be above or below the monopoly price. The bottom of the cycle is at
marginal cost.

lower equal-price market share) has a greater incentive to undercut from equal prices. That is, the small firm leads the large firm down the cycle (behavioral prediction (2)). Conversely, for reasonable parameter values, the large firm is more likely than the small firm to increase price back to the top of the cycle. Noel [2004] further argues that coordination problems make large firms (who control the price for many stations) more natural and effective leaders in price relenting (behavioral prediction (3)). I test these predictions against the data.

It is important to note that Edgeworth Cycles are not restricted to homogeneous Bertrand. Noel [2004] simulates the model above with fluctuating marginal costs and a variety of demand functions D^i, including spatially differentiated markets. The author shows that Edgeworth Cycles are an equilibrium in such markets provided the differentiation is relatively small. The nature of the cycles is similar to that of the homogenous case and in particular the structural and behavioral predictions continue to hold.[2]

Finally, the assumption of alternating moves, on which the theoretical cycles depend, appears to be consistent with industry practice in gasoline markets. Discussions with regional managers suggest firms monitor

---

[2] The structural prediction and behavioral predictions (1) and (2) are robust to starting values used in the simulations. Behavioral prediction (3) can vary depending on starting values, both in the homogeneous and differentiated cases. To allow for greater coordination ability by the large firm, one simply starts off their $V^i(p)$ at lower values for low p. Doing so supports behavioral prediction (3).

competitor prices (easily visible on large billboards) periodically and adjust prices in response. I note that although search costs and menu costs are small, they are positive and determine the frequency of search and price change. Throughout, I take this period to be exogenous and the same for both types of firms. I return to the possibility that different menu costs or search costs may be responsible for generating the cycles when I discuss competing hypotheses below. Lastly, because each firm changes prices after observing its competitors, alternating moves also appears a reasonable description of behavior.

While articles on retail gasoline competition are many,[3] few papers have specifically addressed asymmetric price cycles of this nature. For the United States, Allvine & Patterson [1974] and Castanias & Johnson [1993] note the Edgeworth-like appearance of the cycles in Los Angeles from 1968 to 1972, and present summary statistics on price changes. Eckert [2002] shows how asymmetric cycles similar to Edgeworth Cycles can lead to a finding that price increases are passed through to retail prices more quickly than decreases using weekly data from Windsor, Canada.

Two papers directly examine the impact of small independents on asymmetric prices in testing for Edgeworth Cycles. Using national data for Canada, Eckert [2003] motivates his theoretical model (described above) with interesting correlations between overall price rigidity and year-end concentration ratios for 19 cities and 6 years. The author finds more rigid prices where concentration ratios are higher.

Noel [2007] explicitly models three distinct pricing patterns in Canadian markets – cycles, sticky pricing, and cost-based pricing in 19 cities over 11 years. The author finds price cycles are more prevalent with more small firms, sticky prices less prevalent, and the cycles have shorter periods, greater amplitudes, and are *less* asymmetric. These relationships are consistent with the theories of Edgeworth Cycles.

As mentioned, the two limitations of these articles are the weekly frequency and lack of station-specificity in the data. First, cycles (and fast cycles especially) can be partially obscured. For example, if the observation for that week is collected in the middle of a marketwide relent, the recorded price will be an average of some stations that have relented and others that have not. The duration of the relenting phase will be measured at two weeks (contrary to the structural prediction of the theory). Since undercutting just before and just after the relenting phase are also missed, the measured amplitude is underestimated and asymmetry can be difficult to detect.

---

[3] Articles on retail gasoline competition include: testing oligopoly models of competition and episodic price wars (Slade [1987], Slade [1992]), wholesale-retail passthrough (Borenstein, Cameron & Gilbert [1997], Godby *et. al.* [2000], and many others), mergers (Hastings & Gilbert [2005]), collusion (Borenstein & Shepard [1996]), and multiproduct station pricing (Shepard [1991]).

The second limitation of these articles is that one cannot directly observe the pricing behavior of individual small and large firms which is needed to test the behavioral predictions of the theory. The current dataset, however, is twelve-hourly and station specific, and permits tests of both the structural and behavioral hypotheses of Edgeworth Cycles.[4]

III.   EMPIRICAL FRAMEWORK

For a particular station, two possible pricing regimes are clearly suggested by both the theory and the data:

1. the relenting phase (regime 'R'), and
2. the undercutting phase (regime 'U')

with discrete switching between the two.

The nature of the theoretical Edgeworth Cycles is that the regimes for a particular station are correlated over time. Undercutting phases tend to persist for many consecutive periods while relenting phases tend to last a single period. The current regime thus carries information about the likelihood of the regime in the following period. Therefore I model firm behavior using a two-regime Markov switching regression framework. (A regular switching model does not have this memory feature.)

Also, a latent regime switching framework is appropriate since the true underlying regime at a point in time is unobservable. 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 on the estimated switching probabilities and past play.

The cycle is likely clean enough in the new dataset 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 might even attempt to eyeball the data. However, the Markov switching regression framework is preferable for several reasons. First, it is more general and I show how it can be used to analyze cycle characteristics with data that is not so clean. 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

---

[4] Although the cycles in Toronto are difficult to detect in weekly data (and their characteristics obscured), Noel [2007] is able to find cycles in 84% of weeks overall (and in more than 98% of weeks in five of the last seven years in the data) using weekly data for Toronto from 1989 to 1999.

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.

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

$$(4) \qquad \Delta RETAIL_{st} = \begin{cases} X_{st}^{i'}\beta^i + \varepsilon_{st}^i & \text{with prob. } 1 - \gamma_{st}^i \\ 0 & \text{with prob. } \gamma_{st}^i \end{cases}$$

where $\Delta RETAIL_{st} = RETAIL_{st} - RETAIL_{s,t-1}$ and $RETAIL_{st}$ is the retail price, $(X_{st}^i)'$ 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 $\alpha^i = E(\Delta RETAIL_{st} | X_{st}^i)$.[5] Regimes are station specific so, in principle, each station can follow a cycle of its own.

Because menu costs and monitoring costs are not exactly zero, a period t is of positive and finite length. Moreover, the 'true' length of a period t as determined by gasoline stations is unlikely to be identical to the length of a period chosen by the econometrician when collecting data (in this case, twelve hours). The true length of a period t may even differ across stations. If the time between datapoints is sufficiently short, one will necessarily observe some zero price changes from one data point to the next even if firms were undercutting every 'true' period. Eckert [2003] and Noel [2004] further show that asymmetric firms may price match instead of undercut in response to certain prices, producing more zero price changes. I include a mass point in each regime at zero to account for this. Separating the zeros from the nonzeros allows me to analyze both the actual size of undercuts when they do occur as well as unconditional expected price changes each period.

The regime specifications are built identically and no restrictions are placed on the sign of the price change for inclusion in a given regime. I simply name the regime in which I find prices to rise quickly the 'relenting phase', and the other the 'undercutting phase'. Particulars of each within-regime specification are discussed together with results in later sections.

There are four Markov switching probabilities in total. Let I_{st} be equal to 'R' and 'U' when station s at time t is in the relenting phase regime and the undercutting phase regime respectively. Then the probability that a station switches from regime i in period $t-1$ to regime 'R' in period t is given by the

---

[5] 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.

logit form:

$$\lambda_{st}^{iR} = \Pr(I_{st} = `R'|I_{s,t-1} = i, W_{st}^i)$$

(5)
$$= \frac{\exp(W_{st}^i \theta^i)}{1 + \exp(W_{st}^i \theta^i)}, \quad i = R, U$$

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. Particular specifications discussed below.

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:

(6)
$$Pr(J_{st}^i = 1|I_{st} = i, V_{st}^i) = \gamma_{st}^i = \frac{\exp(V_{st}^i \zeta^i)}{1 + \exp(V_{st}^i \zeta^i)}$$

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.

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. The switching probabilities are estimated by joint non-linear transformations of the core parameters. The switching probabilities and the within-regime estimates are then used to construct the structural characteristics of the cycle such as amplitude, period, and asymmetry. The appendix outlines these derivations in more detail. Standard errors on the constructed variables are calculated by the multivariate delta method.

IV. DATA

I collect and use a new dataset of twelve-hourly retail prices for the same 22 service stations along an assortment of major city routes in central and eastern Toronto over 131 consecutive 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 (integrated into wholesaling and retailing), nine by independents.[6] Twelve

---

[6] Majors are defined as those that are integrated into refining and retailing, independents are retailers only. Major firms generally have a much larger retail presence than independents. Although majors can choose to lease some stations to private dealers, in urban areas and for all
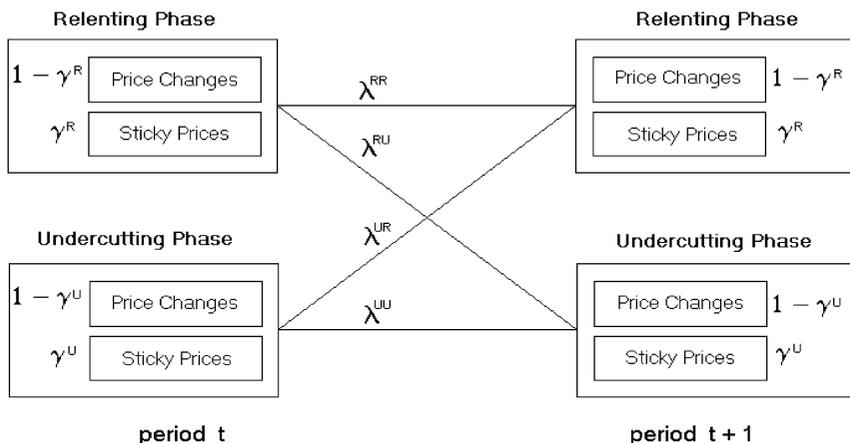
Figure 3
Schematic Overview of Regimes and Switching Probabilities
Each large box represents a regime (relenting or undercutting) and each small box represents a subregime (price changes or sticky prices). The switching probability out of regime i in period t into regime j in period $t + 1$ is given by λ^{ij}. The probability of sticky prices conditional on regime i is given by γ^i.

firms are represented in total including all major national and regional firms. Figure 4 shows a map of all gasoline stations in central and eastern Toronto. The sample stations, spread out over 17 miles, are marked by dark squares.

Retail prices, $RETAIL_{st}$, are for regular unleaded, 87 octane, self-serve gasoline, in Canadian cents per liter (cpl). The descriptive specifications of Section V use after-tax prices (since firms compete on these); the behavioral specifications of Section VI 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.[7] There can be small discounts from this listed price but such discounts are not tied to movements in the retail price. Although only independents buy at rack, the rack price is appropriate since it represents the wholesaler's opportunity cost of wholesale gasoline sold to dealers. Because of readily available U.S. sources of wholesale gasoline, the rack price can be reasonably modeled as exogenous to retail price setting (Hendricks [1996]).

---

major-branded stations in the sample, the head office controls prices. Hence, definitions of 'large' and 'small' are meaningful.

[7] 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.
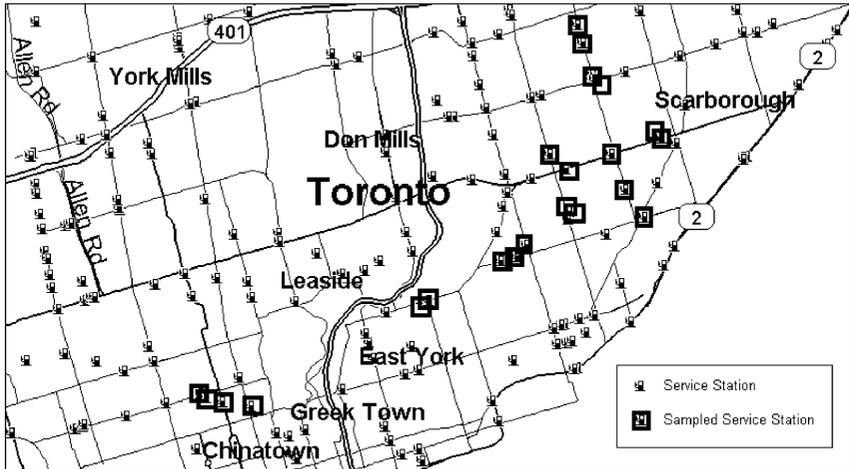
Figure 4
Service Stations and Sampled Service Stations

TABLE I
SUMMARY STATISTICS

	Mean	Std. Dev.	Minimum	Maximum
RETAIL (before tax)	43.09	4.44	32.2	51.8
RETAIL (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

In Canadian cents per liter. Note $POSITION_t = RETAIL_{t-1} - RACK_t$.

Ancillary data such as firm and station characteristics, source of price control and timing of inventory deliveries were self-collected.

Summary statistics for rack and retail prices are shown in Table I. 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.

V.   DESCRIPTION OF THE CYCLE

The main results of this paper are presented in Section VI, when I test the behaviors of small and large firms against predictions of the theory. In this section, I test the structural prediction of the theory and briefly describe the anatomy of a typical cycle – that is, its average amplitude, period, and asymmetry. This is done using a 'summary statistics' specification (specification (1)) in which the within-regime price changes (α^i), switching probabilities (λ^{ij}) and probabilities of sticky pricing conditional on being in regime i, (γ^i), are assumed constant. That is, each X^i, W^i, and V^i are vectors of

TABLE II
WITHIN-REGIME RESULTS AND SWITCHING PROBABILITIES

	(1)	(2)	
	All Firms	Major Firms	Independents
Relenting Phase Estimates (dependent variable: $\Delta RETAIL_{st}$)			
α^R (expected price change)	5.576 (0.083)	5.782 (0.101)	5.192 (0.142)
σ_R (standard deviation of price change)	1.650 (0.066)	1.610 (0.085)	1.655 (0.100)
γ^R (fraction sticky prices)	4.0E-5 (4.6E-4)	4.1E-5 (4.8E-4)	3.8E-5 (4.9E-4)
Undercutting Phase Estimates (dependent variable: $\Delta RETAIL_{st}$)			
α^U (expected price change)	− 0.751 (0.008)	− 0.767 (0.010)	− 0.720 (0.015)
σ_U (standard deviation of price change)	0.459 (0.009)	0.467 (0.012)	0.441 (0.012)
γ^U (fraction of sticky prices)	0.429 (0.007)	0.422 (0.008)	0.441 (0.014)
Switching Probabilities			
λ^{RR} (relenting → relenting)	0.008 (0.004)	0.008 (0.005)	0.007 (0.007)
λ^{RU} (relenting → undercutting)	0.992 (0.004)	0.992 (0.005)	0.993 (0.007)
λ^{UR} (undercutting → relenting)	0.078 (0.001)	0.078 (0.002)	0.078 (0.003)
λ^{UU} (undercutting → undercutting)	0.921 (0.001)	0.921 (0.002)	0.921 (0.003)

Specification (1) does not include firm type dummies. Specification (2) includes firm type dummies in the α^i, λ^{ij}, and, γ^i. Standard errors in parentheses calculated by delta method.

TABLE III
CYCLE CHARACTERISTICS

	(1)	(2)	
	All Firms	Major Firms	Independents
Relating Phase Duration	1.008 (0.004)	1.008 (0.007)	1.007 (0.007)
Undercutting Phase Duration	12.780 (0.291)	12.779 (0.358)	12.784 (0.484)
Cycle period	13.788 (0.291)	13.787 (0.358)	13.792 (0.485)
Asymmetry	12.680 (0.291)	12.677 (0.363)	12.687 (0.491)
Cycle Amplitude	5.619 (0.082)	5.828 (0.098)	5.232 (0.144)

Durations and period in terms of half-day periods, amplitude in cents per liter, measure of asymmetry is unit free. Standard errors in parentheses calculated by the delta method.

ones. Specification (2) repeats the analysis after including dummies for station type (independent and major).

In Table II, I review the within-regime regression results and switching probabilities estimates. These are used to derive the typical structural characteristics of the cycle, as described in the appendix, and reported in Table III.

The evidence supports the structural prediction of the theory that the relenting phase of a given station is complete in a single period followed by a sequence of small consecutive undercuts. The average relenting phase lasts 1.01 half-days and the undercutting phase lasts 12.78 half-days. The expected period of the cycle is therefore 13.78 half-days, or about a week.[8]

---

[8] Possible day-of-the-week effects discussed in Section VII.

I find the expected amplitude of the cycles is 5.61 cpl, 13% of the average ex-tax price, 170% of the average markup, 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 U.S. cents per gallon) were common in the sample. Finally, the cycle is extremely asymmetric. Using the ratio of the undercutting phase duration to the relenting phase duration as a measure of asymmetry, the point estimate of 12.68 is highly significant.

In specification (2), I find virtually identical cycle periods and asymmetries across types of stations, while amplitudes differ. This does not show synchronicity, but suggests a potentially strong interdependence between majors and independents.

VI. SMALL VS. LARGE FIRMS

The theory of Edgeworth Cycles makes specific behavioral predictions about how large and small firms interact and how their pricing behaviors differentially evolve over the path of the cycle. In particular: (1) reactions should be fast so that cycles across stations appear highly synchronous, (2) small firms should lead prices downward and (3) large firms should lead prices upward. The high-frequency, station-specific data used in this study allows a clean test of these behavioral predictions of the model.

I allow for changing behavior along the path of the cycle using two key variables: *POSITION* and *FOLLOW*, described below. Also, because I want to test for differential effects by firm size (major or independent), I interact each of these variables by firm size where they enter the model.

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*.[9] This is done by including *POSITION* in the X^R, X^U, and V^U matrices.

Changes in *POSITION* 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 (λ^{UR} via W^U). I do not include it in the switching probabilities 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 of *POSITION*,

---

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

MICHAEL D. NOEL

TABLE IV
EFFECTS OF CYCLE POSITION

	(3)		
	Major Firms		Independents
$\frac{\partial \lambda^{UR}}{\partial POSITION}$	− 0.065 (0.006)		− 0.051 (0.006)

Switching Probabilities at various values of POSITION			
POSITION =	3.31 mean value overall	1.53 mean value bottom of cycle	0	− 3.24 minimum value
λ^{UR} (undercutting → relating)	0.046 (0.002)	0.110 (0.006)	0.272 (0.022)	0.803 (0.038)
λ^{UU} (undercutting → undercutting)	0.953 (0.002)	0.889 (0.006)	0.727 (0.022)	0.197 (0.038)

Specification (3) adds to specification (2) by allowing *POSITION* to impact the within-regime expected price changes (α^R and α^U), the switching probabilities out of the undercutting phase (λ^{UR} and λ^{UU}), and the probability of sticky pricing in the undercutting phase (γ^U). The reported derivative $\partial\lambda^{UR}/\partial POSITION$ (The change in the probability of switching from undercutting to relenting with respect to *POSITION*) is evaluated at *POSITION* = 1.53, its mean value at the bottom of the cycle. Other parameter estimates are similar to specification (2) and not reported. Standard errors in parentheses calculated by the delta method.

see Table IV. The examples are based on a specification identical to specification (2) but that includes *POSITION* in X^R, X^U, V^U, and W^U. The example shown is for major firms although that for independents is similar.[10] As seen in Table IV, 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 in 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 is easily constructed.[11] Once all have relented, *FOLLOW* is set back to zero for every station. Since I 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

[10] The partial derivative of the switching probability from U to R with respect to *POSITION* is calculated by $\frac{\partial \lambda^{UR}}{\partial POSITION} = \theta^U_{POSITION} * \lambda^{UR} * \lambda^{UU}$.

[11] 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.

TABLE V
LEADERS AND FOLLOWERS

	(4)	
	Major Firms	Independents
$\frac{\partial \alpha^R}{\partial FOLLOW}$	− 0.200 (0.104)	0.084 (0.096)
$\frac{\partial \alpha^R}{\partial POSITION}$	− 0.897 (0.027)	− 0.873 (0.061)
$\frac{\partial \alpha^U}{\partial POSITION}$	− 0.042 (0.007)	− 0.034 (0.027)
$\frac{\partial \gamma^U}{\partial POSITION}$	0.036 (0.002)	− 0.034 (0.005)

Switching Probabilities at $POSITION = 1.53$				
	LEADERS		FOLLOWERS	
	Majors	Independents	Majors	Independents
λ^{UR} (undercutting $\rightarrow$ relenting)	0.091 (0.006)	0.026 (0.003)	0.925 (0.045)	0.720 (0.045)
λ^{UU} (undercutting $\rightarrow$ undercutting)	0.919 (0.006)	0.974 (0.003)	0.075 (0.045)	0.280 (0.045)

Switching Probabilities at $POSITION = 0$				
	LEADERS		FOLLOWERS	
	Majors	Independents	Majors	Independents
λ^{UR} (undercutting $\rightarrow$ relenting)	0.235 (0.020)	0.066 (0.011)	0.974 (0.016)	0.868 (0.030)
λ^{UU} (undercutting $\rightarrow$ undercutting)	0.765 (0.020)	0.934 (0.011)	0.026 (0.004)	0.132 (0.030)

Expected price change in regime i (conditional on a positive price change) is α^i. The probability of sticky pricing conditional on regime U is γ^U. The switching probability from i to regime j is λ^{ij}. Standard errors in parentheses calculated by the delta method.

(λ^{UR}) and the expected price change in the subsequent relenting phase (α^R) to depend on $FOLLOW$.

The results, described below, are reported as specification (4) in Table V. In the specification, X^R and W^U include variables $MAJOR$, $FOLLOW$, $POSITION$, $MAJOR^*FOLLOW$, $MAJOR^*POSITION$, X^U and V^U include variables $MAJOR$, $POSITION$, $MAJOR^*POSITION$, and W^R and V^R include only $MAJOR$ and the constant term.

Assume I have already shown behavioral prediction (1) that cycles for each station are highly synchronous and assume now that all 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. A finding that more active undercutting by *smaller* firms occurs near the tops of the cycles would be consistent with the theory.

In the top half of Table V, 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 $FOLLOW$. Each is reported separately for small (independents) and large (major) firms.

Behavioral prediction (2) is borne out by the data. As predicted, small firms are more aggressive near the top of their cycles and more likely to

trigger new rounds of undercutting. Near the top of the cycles, actual price undercutting (of any size) is substantially more prevalent among the small independents while sticky prices are more prevalent with large major branded firms. Nearer to the bottom, the roles reverse and we see that undercuts are more common with large major firms and sticky prices more common with small independents $\left(\partial\gamma_{MAJ}^{U}/\partial POSITION = 0.036, \partial\gamma_{IND}^{U}/\partial POSITION = -0.034\right)$. 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. Large major 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. As independents continue to undercut major firms and each other, however, the majors respond. In fact, just prior to a new round of relenting, the prices at majors are often below those of the independents.

The undercuts when they do occur are slightly larger at the top of the cycle than at the bottom $\left(\partial\alpha_{i}^{U}/\partial POSITION\right)$. For majors, this appears due to the first undercut from the top of the cycle; for independents, the effect is insignificant.

Behavioral prediction (3) states that larger firms have a greater incentive and greater coordinating ability to trigger a new round of relenting phases once markups become low. Behavioral prediction (1) is that reactions by other firms are so fast that the cycles across stations should appear 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 Edgeworth Cycles.

In the bottom half of Table V, 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 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. Note that because reaction times are so fast, I generally identify a leading group of stations rather than a single leading station, even with twelve-hourly data.

Behavioral prediction (3) – that larger firms are more likely than smaller ones to initiate new rounds of relenting phases – is also confirmed by the data. Consider the switching probabilities at *POSITION* = 1.53 and examine 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 better than the 1% level of significance. This evidence shows that large firms are much more likely to initiate price relenting than small firms.

Next consider the *FOLLOWER* columns. Conditional on at least one station's having relented in the 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. This confirms behavioral prediction (1): firms large and small respond extremely quickly to a new relenting phase of another station by relenting themselves, usually within half a day. The price increases across the city appear to be highly synchronous. Had the data been just bi-daily or less frequent, stations would have appeared to be pricing in perfect synchronicity.

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 react more quickly to a relent by a leading station 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 insight gained from the case when *POSITION* reaches zero is the same. These results are consistent with the existence of Edgeworth Cycles in the Toronto retail gasoline market.

Two more results are worth noting. First, the theory predicts that a following firm will set its price just below that of the leader, 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 because independents tend to all be followers even though the fastest ones make it into the leading group.

Finally, I find that a station's price change during a relenting phase will indeed 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 $\left(\partial\alpha^R_{MAJ}/\partial POSITION = -0.91, \partial\alpha^R_{IND}/\partial POSITION = -0.87\right)$, suggesting that an almost standard markup reinstated each time. Since it is less than one, however, amplitudes become slightly smaller when *POSITION* is lower, for example, due to an increase in wholesale prices.

While the theory states that the top-of-the-cycle price may be either above or below the monopoly price, it appears in these markets to be below given typical estimates of aggregate elasticity. It is an interesting question then

why firms would limit themselves to this amplitude rather than attempt greater ones. One possibility is consumers' intertemporal elasticity of substitution. If consumers would respond to any greater amplitudes by investing in learning how to time their purchases to only periods of low prices (at a cost), it would undo any attempt by firms to increase the amplitude even further.

It seems that few consumers have made the investment in learning about the cycle. In an informal poll 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*. The misperception may be because the large, seemingly simultaneous price increases (especially those resulting in all-time high nominal prices during the sample period) receive much negative press, while the small undercuts each day receive no fanfare at all.

The welfare implication is that better informing consumers about the current cycle (or in other ways lowering their intertemporal substitution costs) may reduce amplitudes and prices even further.

Although the sample stations in this study are geographically spread out over 17 miles, to check their representativeness I also periodically sampled 26 stations in other parts of the city during data collection. I confirm that prices at all these stations moved closely together. The median pairwise price difference between any two stations anywhere in the city was under 0.4 cents and a difference of more than three cents occurred in less than one quarter of one per cent of pairwise comparisons.[12] Moreover, *every* sampled station participated in the cycle. I conclude that the results herein are representative of the city as a whole and support the existence of a single market.

The marketwide nature of the cycle may be surprising. Transportation costs imply some spatial differentiation across stations, but the presence of Edgeworth Cycles would suggest that the differentiation between stations is very low.[13] The fact that prices are highly correlated (independent of marginal cost) even across distant stations, argues that stations are well connected to each other by a chain of competing stations in between. As

---

[12] Not including pairwise comparisons where one firm has already relented but the other has not. During the sample period, 3 cents per liter (Canadian) equals approximately 7.5 cents per gallon (U.S.).

[13] High firm-level price elasticity is an important factor. Imperial Oil Ltd. [2001] 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.) If this is true, perhaps additional utility is being gained from paying the lowest price, since savings would only be about a dime on a fillup. My own anecdotal evidence when collecting this data suggests that a difference of 1 cpl at two nearby stations (*very* rare and very brief) has a large impact on consumer choice.

mentioned earlier, major firms may also contribute to the near marketwide synchronicity of the relenting phases by coordinating across its stations.

VII.   COMPETING HYPOTHESES

Another advantage of the twelve-hourly, station-specific data is that one can more clearly distinguish between Edgeworth Cycles and several competing hypotheses that might explain the asymmetry in prices. In this section, I discuss day-of-the-week demand cycles, menu costs, inventories, rack price discounts and covert collusion as possible alternative explanations.

One competing hypothesis for the cycles focuses on fluctuating demand. That the period of a cycle is about a week long suggests that a day-of-the-week demand cycle may be involved. However, this hypothesis is quickly dispelled. First, it is implausible that gasoline demand would follow the exact pattern consistent with the structural prediction of the theory – a large sudden increase in gasoline demand on one day of the week followed by small decreases in demand every subsequent day. Moreover, the price increase occurs on a different day of the week from one week to the next, and cycle periods range from 4 to 10 days in the sample period (on rare occasion two in the same week). This is inconsistent with a day-of-the-week demand pattern.[14] Also, a demand story does not suggest differential behavior by large and small firms in leading relenting and undercutting phases, as occurs in the data.

It is sensible, though, that varying weekly demand may fine tune the exact timing of the relenting phase in a cycle that would be roughly a week in length anyway. To check this, I performed theoretical simulations in which demand was allowed to fluctuate. Demand could be either high or low in each period (with equal probability) and the active firm learns current demand just prior to setting price. The results show that firms are more likely to relent in the low demand period when the cost of relenting is relatively lower. There is some supporting evidence in the data that relents are more likely to be earlier than later in the week.[15] Quantity data would be needed to say more.

However, I do not find evidence of the often claimed 'long weekend effect' – that is, firms raising prices higher specifically for the long weekend.

---

[14] Noel [2007] shows that in other cities, cycles have periods of several weeks or even several months.
   [15] The probabilities for the *first* firm (generally a major) to relent are: Monday 18%, Tuesday 32%, Wednesday 22%, Thursday 18%, Friday 10%, weekends 0%. Industry activity is low on weekends, which may also explain why Friday appears a poor choice to attempt to trigger a marketwide relent. Using all relents, the percentages are Monday 22%, Tuesday 37%, Wednesday 21%, Thursday 13%, Friday 6%, weekends 1%. The empirical specifications do not include dummy variables for day-of-the-week. Adding early/late week dummies does not impact the previous results.

This is in contrast to a government study which claims to have found one and cites it as evidence of non-competitive behavior.[16] The relenting phases that occur in the week prior to the long weekend are not exceptional, and (although the number of long weekends is small), do not appear to be any later in the week in general.

Differences in menu costs or monitoring costs may be suggested as a second alternative explanation for the differential behaviors found along the cycle. The fact that majors change their prices slightly (albeit insignificantly) more often ($1 - \gamma^U_{MAJ} = 0.58$, $1 - \gamma^U_{IND} = 0.56$) may be suggestive of lower menu and monitoring costs for large firms. Differences in these costs themselves cannot create such a tall and structurally asymmetric cycle. However, even if the cycle were due to some other reason, such as a demand cycle, cost differences cannot explain differential firm behavior. If large firms had lower costs, one would expect them to adjust more quickly in both the upward and downward direction, not only in the upward direction as I find.

A third explanation for the cycle is the depletion of the inventories in the underground tanks at retail stations. For simplicity, assume an exogenous delivery schedule and an effort by stations to exactly deplete inventory prior to the next delivery.[17] Then if a station's sales are lower than expected, inventories do not as fall much as expected, and its price decreases in the next period. However, to decrease repeatedly as in an Edgeworth Cycle, stations would have to repeatedly overestimate its sales in every single period including the period when it chooses its relenting price. If, on the other hand, stations underestimate sales in every single period, one would observe the opposite asymmetry of what I find. Worse than myopia, station managers would need to be perpetually unaware of both the future and the past and systematically err in a precise way.

In practice, gasoline inventories at stations are not scarce and the shadow price of any capacity constraint should be low. Delivery schedules are endogenously set to each station's requirements and extra supply can be readily obtained if needed. Changes in the costs or benefits of holding excess inventory cannot explain a four fold change in markups over a the course of a week, let alone its asymmetry.

Nor is the inventory story consistent with behavioral prediction (1) that reactions are fast and cycles across stations appear highly synchronous. Since deliveries occur on different days for different stations (and is typically less frequent than a week), one would expect any 'cycles' under an inventory

---

[16] Government of Canada [1998] in its review of the downstream industry expressed concern over the synchronicity and volatility of retail gasoline prices. They consider the industry 'tacitly collusive' and postulate a single price leader who moves prices both higher and lower.

[17] Stations may hold some additional inventory for its option value in the event of a positive demand shock prior to the next delivery date. Because station demand is relatively easy to forecast ten days into the future, this is likely to be a small percentage of capacity.

story to be longer and largely independent across stations. The data shows they are not.

A fourth 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. However, rack price discounts are much smaller than the amplitude of the cycle (< 1 cpl versus 5.6 cpl). While they vary by volume purchased, they do not vary over time as required for a cycle. Wholesale supplies are also bought less frequently than the cycle period and they are bought at different times by different stations. (It would also seem strange that a wholesaler would symmetrically change its posted rack prices over time while asymmetrically adjusting any discounts.) Again, this hypothesis cannot explain the differential large and small firm behaviors along the cycle. I note that had there been a rack price cycle instead, it would be just as interesting.

Fifth and finally, consumer groups often cite covert collusion as the cause for what is claimed to be synchronous price movements. This has led to periodic federal investigations in search for evidence of collusion, to date without success.[18] Because the folk theorem teaches that there is an infinite number of possible equilibria in supergames with sufficiently high discount factors, collusion can never be fully ruled out. However, it is extremely unlikely that the cyclical path of prices we observe (which happen to well resemble Edgeworth Cycles) would be the choice of firms colluding under supergame strategies. The most effective collusion strategies in practice are those that are simple to reach, monitor, and punish. Constant price or constant markup rules are examples. These simple strategies involve a minimum of explicit communication and reduce the risk that firms will draw suspicion from antitrust authorities. In contrast, setting up and policing a complicated system of differentially and fast moving prices among hundreds of stations would be very difficult and require plenty of explicit communication. It is also unnecessary. As shown above, there is no evidence of a cost or demand cycle that would suggest any benefit from attempting a complicated cyclical equilibrium instead of a simple one. Moreover, since complaints are often triggered by the large (25%) market price increases in the relenting phase, the cycle would seem a particularly peculiar choice for secretly colluding firms.

Perhaps the leading candidate of all the folk theorem strategies is that firms select and follow a price leader. The price leader would be free to adjust prices as market conditions dictate and other firms are then required to follow. However, the data shows there is no one price leader in the data. One of several different firms may lead prices back to the top of the cycle each

---

[18] I note that the fineness of the data in this study shows the price increases are in fact not perfectly synchronous, but rather a sequence of fast reactions. With data just bi-daily or less frequent data, they would have only appeared to be perfectly simultaneous.

time. When undercutting, many different individual firms lead prices lower and the price ranking of stations changes frequently and unsystematically along the path. These are inconsistent with a simple organizational structure based on a price leader. Even had there been a single price leader, again, the complicated cyclical price pattern we observe would have be a peculiar choice, given no underlying cost or demand cycle to motivate it.

VIII. CONCLUSION

In this paper, I present a new dataset to examine pricing dynamics in the Toronto retail gasoline market. I find evidence consistent with the presence of Edgeworth Cycles, a theoretical construct seemingly implausible in real world practice. The asymmetric shape of the empirical cycle is clear. Consistent with the theory, I find that larger firms are more likely than smaller firms to initiate new rounds of relenting phases and the opposite is true for undercutting phases. The magnitude of relenting phase price increase is 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 cycles also appear highly synchronous across stations. These results are inconsistent with competing explanations for the cycle such as covert collusion and inventory or demand cycles.

The result is interesting in a number of ways. Unlike traditional, periodic price wars which are often seen to facilitate collusion, the 'price wars' here are not punishments triggered by colluding firms. Rather, the competitive outcome involves prices that fall repeatedly and to some extent predictably. Elastic consumers can achieve lower average prices by investing in learning the cycle process and timing purchases accordingly. Coasian dynamics may result to lower prices more broadly. It also shows a leading role for small firms in triggering the undercutting phases that lower prices.

The results contradict the assumption of a single long run steady state price made in most papers of gasoline pricing dynamics in these markets. Although Edgeworth Cycles do not currently appear in U.S. gasoline markets, empirical researchers working in markets with similar characteristics need to consider the potential for this sort of pricing dynamics in their estimation. Where else Edgeworth Cycles might appear remains to be seen.

APPENDIX

There are two top-level regimes: $I_{st} = $ 'R', 'U'. Each is subdivided into subregimes $J_{mt} = \{0,1\}$ for non-sticky and sticky pricing respectively. The closed form log likelihood function for the Markov switching model is computationally intractable and so is computed by means of a recurrence relation, as described by Cosslett & Lee [1985]. Let

$$(7) \qquad Q_{st}(I_{st}) = \sum_{I_{st-1}=R,U} g^{I_{st}}(\varepsilon_{st}^{I_{st}}|X_{st}^{I_{st}}, V_{st}^{I_{st}}) * \Pr(I_{st}|I_{s,t-1}, W_{st}^{I_{st}}) * Q_{s,t-1}(I_{s,t-1})$$

where

$$
\begin{aligned}
(8) \quad g^{I_{st}}(\varepsilon_{st}^{I_{st}}|X_{st}^{I_{st}}, V_{st}^{I_{st}}) =& \Pr(J_{st}=0|V_{st}^{I_{st}}) * \phi(\varepsilon_{st}^{I_{st}}|X_{st}^{I_{st}}) \\
&+ \Pr(J_{st}=1|V_{st}^{I_{st}}) * D(p_{st}-p_{s,t-1})
\end{aligned}
$$

and where ϕ is the normal pdf, $D(x)$ is an indicator variable equal to 1 if $x=0$, and $Q_{s1}(I_{s1})$ are chosen starting within-regime probabilities. Note that $\Pr(I_{st}=j|I_{s,t-1}=i)$ is called λ^{ij} and $Pr(J_{st}=1|V_{st}^{I_{st}})$ is called γ^i in the text. Then the likelihood function is computed by

$$
(9) \quad L = \sum_{s=1}^{S}\sum_{t=1}^{T}\ln\left(\sum_{I_{st}=R,U} Q_{st}(I_{st})\right)
$$

Results are not sensitive to starting values and, given the crispness of the data, converged easily.

The structural characteristics of the cycles are calculated directly from the switching probabilities and the within-regime parameters as follows:

$$
(10) \quad E(duration\,of\,regime\,i) = \frac{1}{1-\lambda^{ii}}
$$

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

$$
(12) \quad E(amplitude) = \frac{(1-\gamma^R)\alpha^R}{1-\lambda^{RR}}\left(\text{or}\,\frac{-(1-\gamma^U)\alpha^U}{1-\lambda^{UU}}\right)
$$

$$
(13) \quad E(asymmetry) = \frac{1-\lambda^{RR}}{1-\lambda^{UU}}\left(\text{or}\,\frac{-(1-\gamma^R)\alpha^R}{(1-\gamma^U)\alpha^U}\right)
$$

In the text, the first equation of each pair is used for the amplitude and asymmetry measures.

REFERENCES

Allvine, F. and Patterson, J., 1974, *Highway Robbery, An Analysis of the Gasoline Crisis*, (Indiana University Press, Bloomington, Indiana). pp. 192–197, p. 243.

Borenstein, S. and Shepard, A., 1996, 'Dynamic Pricing in Retail Gasoline Markets,' *RAND Journal of Economics*, 27, pp. 429–451.

Borenstein, S.; Cameron, A. C. and Gilbert, R., 1997, 'Do Gasoline Markets Respond Asymmetrically to Crude Oil Price Changes?,' *Quarterly Journal of Economics*, 112, pp. 305–339.

Castanias, R. and Johnson, H., 1993, 'Gas Wars: Retail Gasoline Price Fluctuations,' *Review of Economics and Statistics*, 75, pp. 171–174.

Cosslett, S. and Lee, L., 1985, 'Serial Correlation in Latent Variable Models,' *Journal of Econometrics*, 27, pp. 79–97.

Eckert, A., 2002, 'Retail Price Cycles and Response Asymmetry,' *Canadian Journal of Economics*, 35, pp. 52–77.

Eckert, A., 2003, 'Retail Price Cycles and Presence of Small Firms,' *International Journal of Industrial Organization*, 21, pp. 151–170.

Edgeworth, F. Y., 1925, 'The Pure Theory of Monopoly,' in *Papers Relating to Political Economy*, Volume 1 (MacMillan, London). pp. 111–142.

Ellison, G., 1994, 'Theories of Cartel Stability and the Joint Executive Committee,' *RAND Journal of Economics*, 25, pp. 37–57.

Godby, R. A., Lintner, A., Stengos, T., and Wandschneider, B., 2000, 'Testing for Asymmetric Pricing in the Canadian Retail Gasoline Market,' *Energy Economics*, 18, pp. 349–368.

Government of Canada 1998. 'Report of the Liberal Committee on Gasoline Pricing in Canada,' (Ottawa), pp. 17–20.

Hendricks, K., 1996. 'Analysis and Opinion on Retail Gas Inquiry,' report prepared for Industry Canada, pp. 1–4.

Hastings, J. and Gilbert, R., 2005, 'Market Power, Vertical Integration, and the Wholesale Price of Gasoline,' *Journal of Industrial Economics*, 53, pp. 469–492.

Imperial Oil Limited 2001, *http://www.esso.ca/news/issues/mn_price*, accessed July 1, 2001.

Maskin, E. and Tirole, J., 1988, 'A Theory of Dynamic Oligopoly II: Price Competition, Kinked Demand Curves and Edgeworth Cycles,' *Econometrica*, 56, pp. 571–599.

Noel, M., 2004, 'Edgeworth Price Cycles and Focal Prices: Computational Dynamic Markov Equilibria,' UCSD Working Paper, 2004–13, pp. 1–32.

Noel, M., 2007, 'Edgeworth Price Cycles, Cost-based Pricing and Sticky Pricing in Retail Gasoline Retail Markets,' *Review of Economics and Statistics*, forthcoming.

Shepard, A., 1991, 'Price Discrimination and Retail Configuration,' *Journal of Political Economy*, 99, pp. 30–53.

Slade, M., 1987, 'Interfirm Rivalry in a Repeated Game: An Empirical Test of Tacit Collusion,' *Journal of Industrial Economics*, 35, pp. 499–516.

Slade, M., 1992, 'Vancouver's Gasoline Price War: An Empirical Exercise in Uncovering Supergame Strategies,' *Review of Economic Studies*, 59, p. 257–276.