

UC San Diego

Recent Work

Title

Do Gasoline Prices Resond Asymmetrically to Cost Shocks? The Confounding Effect of Edgeworth Cycles

Permalink

<https://escholarship.org/uc/item/7cx539sc>

Author

Noel, Michael

Publication Date

2007-06-05

2007-04

UNIVERSITY OF CALIFORNIA, SAN DIEGO

DEPARTMENT OF ECONOMICS

Do Gasoline Prices Respond Asymmetrically to Cost Shocks? The
Confounding Effect of Edgeworth Cycles

By

Michael Noel
Department of Economics
University of California, San Diego

DISCUSSION PAPER 2007-04

June 2007

Do Gasoline Prices Respond Asymmetrically to Cost Shocks? The Confounding Effect of Edgeworth Cycles

Michael Noel
University of California, San Diego

June 5, 2007

Abstract

Asymmetric price cycles which look similar to Edgeworth Cycles are appearing in increasingly many retail gasoline markets in the U.S. and worldwide. The cycles can give the appearance of asymmetric price *responses* to cost shocks under traditional methodologies. This article shows how to remove the confounding effect of the cycles and test for any *true* underlying asymmetry in price responses. Designing the correct counterfactual is key. The methodology is demonstrated for one strongly cycling market and some asymmetry to cost shocks is found. Covert collusion is unlikely, but the ability to coordinate cyclical price increases may play a role. Consumers can still reduce expenditures on gasoline up to 7.7% with simple timing rules of thumb.

JEL Classification L13, L41, L81

1 Introduction

This article addresses an increasingly important problem when testing for asymmetric price responses to cost shocks in retail gasoline markets. The source of the problem lies at the intersection of two branches of the literature – the asymmetric passthrough literature and the Edgeworth Cycles literature – which I first need to preview.

Asymmetric passthrough – the idea that positive cost shocks are passed through to prices more quickly than negative ones – has generated much interest by academics, consumers, and policymakers alike. If downstream prices rise faster after an increase in upstream prices than they fall after a decrease, there are clear welfare and policy concerns, especially with market power as a potential cause. A large literature has developed to examine this phenomenon in gasoline markets in the U.S., Canada, and many other countries. A widely cited but controversial work is Borenstein, Cameron, & Gilbert [1997], who find asymmetric passthrough from crude to retail prices in the U.S. They attribute the asymmetry partly to retailer market power and tacit collusion. Others argue for explanations other than collusion (Lewis[2004], Radchenko[2005]) or find no asymmetry at all (Bachmeier & Griffin[2003]).¹

A new strand of the gasoline markets literature that has generated as much recent interest is that of Edgeworth Cycles. With better data now available, a rapid and asymmetric retail price cycle is being detected in many retail gasoline markets around the world. Numerous authors have studied these cycles separately and argue they are the theoretical Edgeworth Cycles discussed by Maskin & Tirole[1988]. In an Edgeworth Cycle, firms undercut one another slowly until prices get close to the wholesale price. At that point, one firm raises price sharply, others follow at once and a new round of undercutting begins. Once considered only a theoretical construct, cycles believed to be Edgeworth Cycles have now been documented in the Canadian (Noel[2007a, 2007b], Eckert[2002,2003]), U.S. (Castanias & Johnson [1993])

¹Other studies using VAR, PAM, or related methodologies include Bacon[1991], Manning[1991] and Reilly & Witt[1999] for the U.K., Godby et al.[2000] and Hendricks[1996] for Canada, Kirshgassner & Kubler for Germany and Duffy-Deno[1996], Karrenbrock[1991] and Balke et al.[1999] for the U.S.

and Australian (Wang[2005]) gasoline markets, in several European countries, and in internet auctions (Zhang & Feng [2005]). I defer to these papers for further evidence that the empirical cycles are Edgeworth Cycles.

Conceptually, the asymmetry in prices caused by Edgeworth Cycles is separate from the asymmetry of price *responses* to wholesale price shocks. For example, cycles generate price asymmetry absent any shocks at all. And response asymmetry, if it exists, does not require cycles. But there exists an interesting interaction between the two. As I will make more precise later, wholesale price shocks can shift cycle peaks forward or backward in time and thereby create the appearance of asymmetric passthrough. This would be true even if every component of the cycle responds to shocks perfectly symmetrically. The detected asymmetric passthrough would then be purely the by-product of the cycle and not the result of the many hypotheses, including collusion or search costs, found in the literature. It turns out that, with one exception, all past studies of asymmetric passthrough – including a few conducted in markets now known to experience rapid cycles – do not allow for self-generating asymmetric prices in their estimation. Where cycles do not really exist, this will be perfectly fine. But where they do (even if they are difficult to see in the data), this creates a bias.

In an important predecessor to this article and the exception to the rule, Eckert[2002] nicely shows that the retail price cycle can, under standard methodologies, give the appearance of asymmetric passthrough. The author attributes the asymmetric response he finds in his study market of Windsor, Canada to the presence of the strong price cycles there.

But as I explain in more detail below, the detected asymmetric passthrough remains a mix of two effects – 1.) the confounding price asymmetry of the Edgeworth Cycle itself and 2.) the true difference in the future price paths across actual and counterfactual worlds that is triggered by the shock. The existence of the first effect does not eliminate the potential of the second and it is the second whose sources are not well understood and which remains of independent policy interest. The two effects combined I call “*apparent*

asymmetric passthrough”. The second effect alone I call “*true* asymmetric passthrough”. Clearly, it would be incorrect to conclude *true* asymmetric passthrough exists when in fact only Edgeworth Cycles are generating all the *apparent* asymmetric passthrough. But it would be equally incorrect to disregard the potential for *true* asymmetric passthrough simply because a cycle is present. Especially with ongoing concerns of collusion as a potential source for *true* asymmetric passthrough, it is important to separate out and analyze any residual *true* component.

This begs the important question that is the focus of the current article. How can researchers design a test for *true* asymmetric passthrough in the presence of self-generating price asymmetries such as Edgeworth Cycles? In this article, I outline a general approach for distinguishing the two sources and then develop a specific econometric model to perform the test in the context of a market that experiences Edgeworth Cycles. This involves nesting a traditional model of asymmetric passthrough inside a model of Edgeworth Cycles.

The deeper point in this analysis extends beyond my particular study market and even beyond the growing number of gasoline markets around the world where Edgeworth Cycles are being detected. In designing asymmetric passthrough studies, researchers must give careful attention to choosing and estimating the correct counterfactual – i.e. what would have occurred in the absence of a wholesale price shock. The assumed counterfactual in virtually all previous studies on asymmetric passthrough is that absent the shock the price in the current period would have been the same as the price in the previous period. The correct counterfactual, however, is what the price *would have been* in the current period in the absence of the wholesale price shock. In markets where prices change asymmetrically independent of wholesale prices, such as in cycling markets, these are not the same thing. The wrong counterfactual generates a finding of *apparent* asymmetric passthrough; the correct counterfactual gives *true* asymmetric passthrough.

The assumption in the literature of a singleton long run steady-state price seemed quite

reasonable given the available data and in many cases will no doubt remain perfectly valid. But the discovery of Edgeworth Cycles in increasingly many gasoline markets around the world represents an important violation of the almost universal assumption in the literature. Cycles have now been widely found in U.S., Canadian, Australian, and numerous European markets, including some where asymmetric passthrough studies have previously been done.² Only Eckert[2002] accounts for the cycles as a source of apparent asymmetric passthrough although he implicitly constrains true asymmetric passthrough to be zero.

I demonstrate my methodology on a high-frequency, station-specific dataset for the city of Toronto, Canada which experiences strong cycles.³ Peeling away the confounding influence of the Edgeworth Cycle, I find that true asymmetric passthrough does exist in the sample market. The effect is statistically significant but very small economically and ends quickly. The dominant source of apparent asymmetry is the cycles. I evaluate hypotheses for the true asymmetric passthrough and the potential for any concern. I also show it is very easy to undo any negative effects. With surprisingly simple rules of thumb for timing purchases, I find that consumers can reduce total gasoline expenditures overall on the order of 3% to 8%. These rules of thumb can be fine tuned for use in other cycling markets in the U.S. and around the world.

I proceed as follows. In section 2, I outline the Edgeworth Cycle theory. In section 3, I contrast true and apparent asymmetric passthrough in more detail. In section 4, I specify my econometric model and in section 5 I describe the data. I present the main results in section 6 and discuss implications in section 7. Section 8 concludes.

²The sample used by Godby et al.[2000] includes many markets that experience strong cycles, including the market studied here.

³I focus on the wholesale-retail transmission since Edgeworth cycles appear only in retail prices and not in upstream markets.

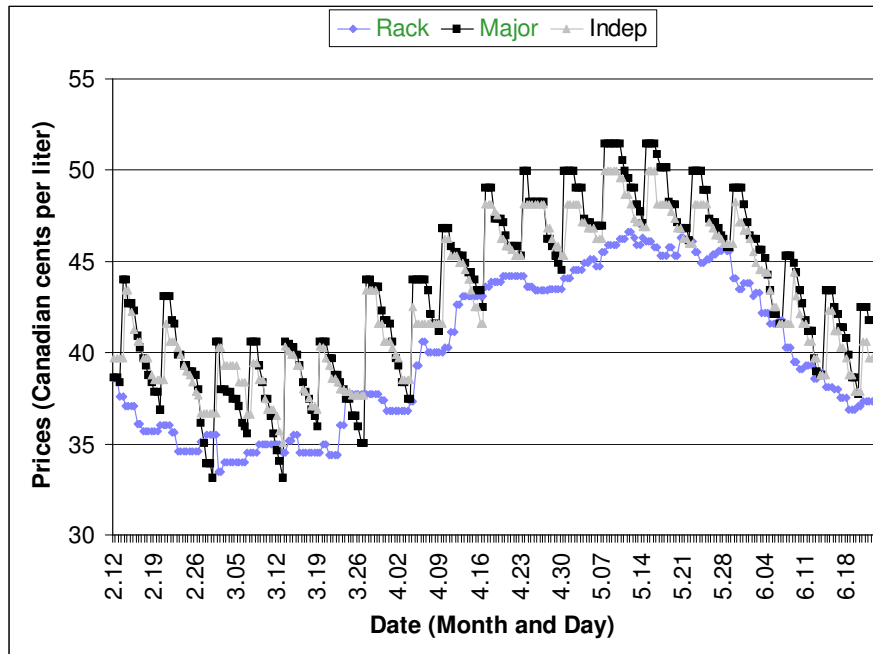


Figure 1: Retail (Major and Independent) and Rack Prices, Feb. 2001

2 Edgeworth Cycles

Asymmetric price cycles that appear to be Edgeworth Cycles have now been documented in a growing number of gasoline markets. The sample market chosen in this study, Toronto, Canada, exhibits a very strong cycle. Figure 1 shows the retail prices at a representative station operated by a major national firm and one operated by an independent firm, along with the wholesale or “rack” price, for a six month period in Toronto. These two stations are on different streets approximately five miles apart. The frequency of the data is twelve hours.

Contrast this to Figure 2, which shows an example path of the market price along under theoretical Edgeworth Cycles. I refer to the reader to Noel[2007a,2007b] for evidence that the empirical cycles in Toronto are in fact generated by an Edgeworth Cycle process.

Maskin & Tirole[1988] derive the theoretical Edgeworth Cycles as a Markov Perfect Equilibrium of a dynamic Bertrand pricing game between two homogeneous goods firms




Figure 2: Market Price in an Edgeworth Cycle

with constant marginal costs.⁴ Each firm sets price alternately: if it is Firm 1's turn to adjust its price in period t , it chooses from a discrete price grid and $p_t^1 = p_{t+1}^1$ and $p_t^2 = p_{t-1}^2$. Each firm earns current period profits of

$$\pi_t^i(p_t^1, p_t^2) = D^i(p_t^1, p_t^2) * (p_t^i - c) \quad (1)$$

where D^i is the Bertrand demand function.

The strategies of each firm are allowed to depend only on the payoff-relevant state in each period which in this case is simply the other firm's price set in the previous period. Let $V^1(p_{t-1}^2)$ be Firm 1's value function when Firm 2 adjusted its price to p_{t-1}^2 in the previous period and Firm 1 adjusts its price in the current period. Let $W^1(p_{s-1}^1)$ be Firm 1's value function when it has set price p_{s-1}^1 in the previous period and Firm 2 is about to adjust its price. V^1 and W^1 can be written as

$$V^1(p_{t-1}^2) = \max_{p_t} [\pi_t^1(p_t, p_{t-1}^2) + \delta W^1(p_t)] \quad (2)$$

⁴See Noel[2004] for extensions of the Maskin&Tirole model to fluctuating marginal costs, differentiation, capacity constraints, and triopoly.

$$W^1(p_{s-1}^1) = \mathbb{E}_{p_s} [\pi_s^1(p_{s-1}^1, p_s) + \delta V^1(p_s)]$$

where δ is the discount factor. Similar equations are found for V^2 and W^2 . The expectation in W^1 is taken with respect to the distribution of the best response function of Firm 2, R^2 .

To choose the best response price, given the current rival price p_{t-1}^2 , Firm 1 maximizes V^1 . Firm 2 acts similarly. The equilibrium strategies are given by R^1 , R^2 , where $(p_t^1)^* = R^1(p_{t-1}^2)$, $(p_t^2)^* = R^2(p_{t-1}^1)$ and p_{t-1}^j is the price chosen by Firm j in period $t - 1$ which remains in effect in period t .⁵ Maskin & Tirole show that Edgeworth Cycles is one of two types of Markov Perfect Equilibria that can emerge in this model (the other being focal prices).

In an Edgeworth Cycle, firms undercut one another's prices repeatedly by small amounts to steal market share (the "undercutting phase"). Once markups fall to zero and undercutting is no longer profitable, a war of attrition ensues until one firm relents and resets its price to a substantially higher level ("the relenting phase"). The new price at the top of the cycle can even be higher than the monopoly price. All other firms quickly follow the price increase and the next round of undercutting begins. This repeats over and over, creating the clear sawtooth pattern in the figures.

Note that the cycle generates strongly asymmetric retail price changes – large price increases and small price decreases – even with constant marginal costs.

3 True vs. Apparent Asymmetric Passthrough

Noel[2004] extends the Maskin & Tirole theoretical model to allow for fluctuating marginal costs. Noel shows, consistent with the basic model, that firms are more likely to relent (i.e. raise prices to the top of the cycle) when price-cost margins become low. But price-cost margins now decline through two mechanisms – by undercutting and by positive shocks to

⁵See Maskin & Tirole (2001) for general properties of MPEs (not simply with alternating moves games).

firms' marginal cost. Noel finds that when margins are already low, a positive cost shock squeezes margins and increases the probability that firms will relent quickly.

This turns out to be a source of the apparent asymmetric passthrough. A very large increase in retail gasoline prices often follows a positive shock to rack prices (i.e. marginal costs), since the cost shock can trigger a new relenting phase and price increases in the relenting phase tend to be large. A negative shock to rack prices, however, would not typically be followed by a large retail price decrease, since it gives more room for firms to undercut and undercuts tend to be small. The result is that cost increases appear to be passed through more quickly than decreases. This is the excellent revelation of Eckert[2002] who studies the cycles in Windsor, Canada and concludes they are a source for at least some part of the apparent asymmetric passthrough there.

But the relevant dimension from a policy point of view remains *true* asymmetric passthrough. As explained below, the Edgeworth Cycle may fully account for apparent asymmetric passthrough, or there may still be significant *true* asymmetric passthrough net of the cycles' influence. I test this. The essence of the test is whether the cycle, component-by-component, responds differentially to wholesale price shocks. For example, do firms hurry the next relenting phase (and hence the next cycle) after a rack price increase more than they delay it after a rack price decrease? If we find the sign of past wholesale price shocks differentially impacting the individual cycle components, we have evidence of true asymmetric passthrough acting on the cycles.

This is because the textbook Edgeworth Cycles is a symmetric object in the passthrough dimension. It is symmetric in the sense that each component responds symmetrically to wholesale price shocks. The probability of switching between the two phases of the cycle and magnitude of price changes conditional on the phase do not depend on the direction of past wholesale shocks. Firm pricing decisions depend on current marginal cost and the current prices of its competitors. How the wholesale price got to where it is (by increasing or




Figure 3: True Rack-Retail Passthrough under Correct Counterfactual

decreasing) is not material. If direction matters, firms are deviating from the normal cyclical pattern to pass through shocks asymmetrically.

The general problem with the traditional methodology applied to cycling markets – or any market with self-generating asymmetric price paths – is that it rests on the wrong counterfactual. It compares the prices after the shock to prices before the shock when we should compare prices after the shock to what they would have been after the shock in the absence of the shock. In markets such as cycling markets, these are not the same thing. The wrong counterfactual yields apparent asymmetric passthrough; the correct counterfactual yields true asymmetric passthrough.

Figure 4 gives a highly stylized graphical exposition of this: Panel A shows the effect of a rack price increase on retail prices and Panel B shows it for a decrease. Imagine that the price jump would have occurred with probability one at time t , but the shock to rack prices triggers it earlier (Panel A) or later (Panel B). The apparent asymmetric passthrough is measured by comparing the observed vertical distance XY in Panel A to the observed vertical distance $Y'Z'$ in Panel B. True asymmetric passthrough after the first period, however, occurs when

the distance XZ in Panel A – which represent the price difference between the real and counterfactual worlds – is not equal to the distance $X'Z'$ in Panel B.⁶

Comparing the real and counterfactual worlds, the test I perform is whether the change in the expected price change (i.e. second difference) following a rack price increase is the same as after a decrease (in absolute values). This in turn depends on changes in the probabilities of each phase of the cycle and changes in the expected price change conditional on the cycle phase.

The figures also make clear the critical bias in apparent asymmetric passthrough. True and apparent passthrough are relatively similar in the case of cost increases (XZ versus XY), but very different in the case of cost decreases ($X'Z'$ versus $Y'Z'$). The large price increases that occur after positive cost shocks are easily observable in the data, but the large price increases that do *not* occur because they are *preempted* by negative cost shocks are not observable.

Note lastly that true actual-counterfactual price spread in the diagrams – i.e. the cumulative response functions – can oscillate over time. For example, in the center of each diagram the spread is large (XZ or $X'Z'$) but at the right of the each diagram, it is temporarily reduced to zero. Since the next actual price jump will occur earlier in Panel A than in Panel B, the spread will become large again, then return to zero, then become large, and so on. This will become important soon.

4 A Nested Asymmetric Response Model

To model this econometrically, I use a latent regime Markov switching regression model. For a particular retail station, two possible pricing regimes are clearly suggested by both the

⁶The diagram is illustrative. Notice that in Panel A, the actual and counterfactual cycles diverge before period t and in Panel B, they separate after line t . When infinitesimal changes in rack price are considered and expected price changes and regimes are probability-weighted, the separation between actual and counterfactual in each diagram occurs right at t . The comparisons in retail prices following positive and negative rack price shocks are then made consistently within each time period $\tau \geq t$.

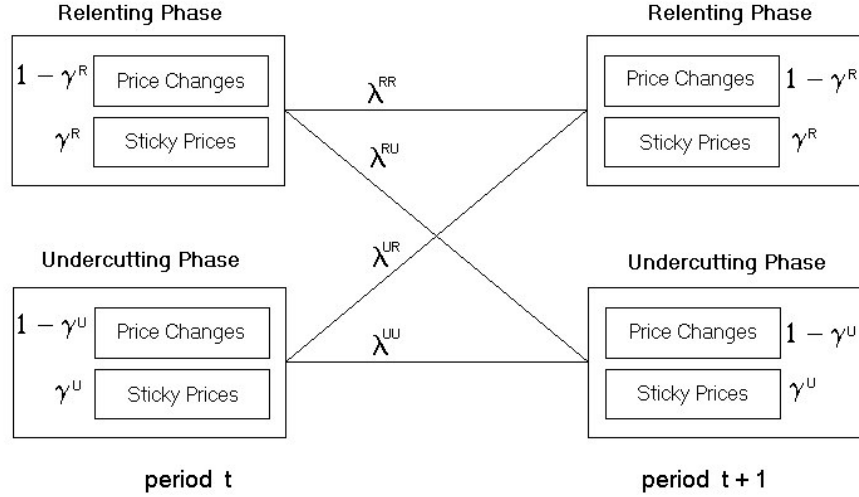


Figure 4: Regimes and Switching Probabilities

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 the Markov 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.

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

$$\Delta RETAIL_{st} = \begin{cases} X_{st}^i \beta^i + \varepsilon_{st}^i & \text{with prob. } 1 - \gamma_{st}^i \\ 0 & \text{with prob. } \gamma_{st}^i \end{cases} \quad (3)$$

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 . Explanatory variables in the final specification include the current rack price $RACK_t$ and previous retail price $RETAIL_{t-1}$ which are relevant for textbook Edgeworth Cycles, and a series of lagged rack price differences that are direction specific, $\Delta RACK_{t-r}^+ = \max(0, \Delta RACK_{t-r})$ and $\Delta RACK_{t-r}^- = \min(0, \Delta RACK_{t-r})$, relevant to testing for true passthrough. Choice of explanatory variables are discussed in more detail below when particular specifications are discussed. Let $\alpha^i = E(\Delta RETAIL_{st} | X_{st}^i)$.⁷ Regimes are station specific so, in principle, each station can follow a cycle of its own.

Because menu costs and monitoring costs are not zero (although they are very low), a period t is of 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. Noel[2004] further shows that 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

⁷Rather 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.

well as overall price changes per period of time.

The regime specifications are built identically and no restrictions are placed on the sign of the retail 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.

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 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 logit form:

$$\begin{aligned}\lambda_{st}^{iR} &= \Pr(I_{st} = \text{“R”} \mid I_{s,t-1} = i, W_{st}^i) \\ &= \frac{\exp(W_{st}^i \theta^i)}{1 + \exp(W_{st}^i \theta^i)}, \quad i = R, U\end{aligned}$$

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. Explanatory variables in the final specification include lagged retail price $RETAIL_{t-1}$, current rack $RACK_{t-1}$ and a series of lagged rack price differences $\Delta RACK_{t-s}^+$ and $\Delta RACK_{t-s}^-$. More on the particulars of each specification 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:

$$\Pr(J_{st}^i = 1 \mid I_{st} = i, V_{st}^i) = \gamma_{st}^i = \frac{\exp(V_{st}^i \zeta^i)}{1 + \exp(V_{st}^i \zeta^i)} \quad (4)$$

where $(V_{st}^i)'$ is a $Q^i \times 1$ vector of explanatory variables and ζ^i is an $Q^i \times 1$ vector of parameters. Again, lagged retail, current rack and a series of lagged rack price differences are included.

Figure 4 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. Standard errors on the constructed variables are calculated by the multivariate delta method.

I then use the parameters to estimate the price paths under three scenarios – 1. with no shock to rack prices, 2. with a positive shock to rack prices, and 3. with a negative shock to rack prices. I perform five hundred simulations of each scenario. Comparing scenario 2 to scenario 1, I get a measure of passthrough of rack price increases and comparing 1 to 3, I get passthrough of rack price decreases. True asymmetric passthrough occurs when the rates of passthrough significantly differ.

5 Data

I use a self-collected dataset of twice-daily 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 12th and June 22nd 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. Twelve firms are represented in total including all major national and regional firms. Figure 5 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



Figure 5: Toronto Gasoline Stations

Canadian cents per liter (cpl), and before taxes. 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.⁸ 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 (Lerner[1996]). 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]).

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.

While I do not reject the null hypothesis that the *RETAIL* and *RACK* series have a

⁸A 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 6: Apparent Asymmetric Passthrough

unit root, of interest in the analysis will be the per period retail price change $\Delta RETAIL = RETAIL_t - RETAIL_{t-1}$ and the difference $RETAIL_{t-1} - RACK_t$. Augmented Dickey Fuller tests on each of these series rejects a unit root in each case (t statistics are -3.60 and -3.81 respectively, compared to the 5% critical value -3.43) in favor of the alternative that each series is stationary. It is still worth noting that the level series $RETAIL$ and $RACK$ are cointegrated. An ADF test using the residuals from an OLS regression of $RETAIL$ on $RACK$ and a trend rejects the null hypothesis of no cointegration (t statistic of 4.11 compared to a 5% critical value of 3.43). Twenty lags were used in all ADF tests.

6 Testing for True Asymmetric Passthrough

Before turning to the main analysis, I examine apparent asymmetric passthrough using a traditional VAR model with an error correction term that is common in the asymmetric passthrough literature. I follow Borenstein et al.[1997] closely and regress first price dif-

ferences on lagged retail and rack prices and lagged retail and rack price differences. The included lag length is 8 and the best response functions are shown out to 15 periods. The coefficients of the underlying regression are difficult to interpret, so in Figure 6 I show the generated cumulative response functions following a positive and negative wholesale price shock. The figure shows price increases are passed through more quickly than price decreases, a common finding in the apparent asymmetric passthrough literature. Typical explanations forwarded by authors for this result in other markets include collusion, inventory asymmetries and search. Eckert[2002] cites Edgeworth Cycles for Windsor.

The reader should be warned however that this result is not robust. Results vary depending on the frequency of the data employed (twice-daily, daily, weekly, etc.), the subsample of data used, and the lag length. This is a common theme in the literature (e.g. Bachmeier & Griffin [2003]). With some subsamples, the reverse asymmetry can even occur.⁹ The lack of robustness in this case is not surprising because of the obvious specification error: retail prices follow a strong asymmetric cycle entirely ignored by the model.

Therefore I use the Markov switching model described earlier to endogenize the cycles. I first perform two preliminary symmetric specifications using this model that will prove helpful later and then proceed to the main exercise of testing for true asymmetric passthrough.

The first specification gives quantitative evidence of the sharp price asymmetry in the cycle. It can be thought of as a “summary statistics” specification of the features of the cycles post-regime categorization. It is estimated by setting each W^i , V^i , and X^i to a column of ones, which allows for a single estimated measure of each relevant characteristic of the cycle. In particular, current or lagged rack prices are not included for this descriptive specification. The results, shown as specification (1) in Tables 2 and 3, confirms the Edgeworth asymmetry clearly seen in Figure 1.

⁹During periods of falling rack prices, there is a spurious correlation between the negative rack price changes and the negative retail price changes from undercutting.

Examining Table 2, the average price change in a relenting phase is 5.57 cpl, whereas the average price change in an undercutting phase is only 0.75 cpl conditional a non-zero change. (Price changes in an undercutting phase are nonzero $1 - \gamma^U = 57\%$ of the time) . It is extremely unlikely for two consecutive relenting phases to occur ($\lambda^{RR} = 0.008$) but undercutting phases are persistent over time ($\lambda^{UU} = 1 - \lambda^{UR} = 1 - 0.078 = 0.922$).

Using the formulas described in Noel[2007b], I calculate the structural characteristics of the typical cycle and report them in Table 3. I find that the duration of a relenting phase is half a day, the duration of an undercutting phase is 12.78 half day periods and the period of cycle is 13.79 periods. The degree of asymmetry – as measured by the ratio of the expected duration of the undercutting phase to the expected duration of the relenting phase – is 12.68 periods and very significantly different from one. The amplitude of the cycle is 5.62 cents, or about 170% of the average *RETAIL* – *RACK* markup. While the asymmetry in retail prices is strong, there is no appreciable asymmetry in the rack price. The evidence points strongly to an asymmetric Edgeworth-like cycle in the Toronto market.

The second and final preliminary specification takes a closer look at the rack-retail transmission mechanism which will prove important later. The theory of Edgeworth Cycles predicts that as *RETAIL*_{*t*-1} decreases or *RACK*_{*t*} increases (i.e. the firm is closer to the bottom of its cycle), the probability that a firm relents will increase (λ^{UR} increases, λ^{UU} decreases.)¹⁰ I thus include these variables in the switching probabilities out of the undercutting phase. Similarly, I include *RETAIL*_{*t*-1} and *RACK*_{*t*} in both within-regime price change equations to allow for within-regime movements to differ depending on the current cycle position. When *RETAIL*_{*t*-1} decreases or *RACK*_{*t*} increases, the magnitude of the price change in each regime is likely to increase. I also include *RETAIL*_{*t*-1} and *RACK*_{*t*} variables into the probability of sticky pricing within the undercutting regime for completeness. Constants and

¹⁰More specifically, in Maskin & Tirole[1988], $\lambda^{UR} = 0$ for *RETAIL*_{*t*-1} > *RACK*_{*t*} and $\lambda^{UR} \in [0, 1]$ for *RETAIL*_{*t*-1} = *RACK*_{*t*}. In Noel[2004] where costs are allowed to fluctuate, λ^{UR} increases more gradually as *RETAIL*_{*t*-1} → *RACK*_{*t*}. Both models find $\lambda^{RR} = 0$. Consistent with these models, two consecutive relenting phases are extremely unlikely in the data, so I do not include *RETAIL*_{*t*-1} or *RACK*_{*t*} in λ^{RR} .

firm type dummies are included in the within regime regressions, in the probability of sticky pricing within an undercutting phase and in the switching probability out of an undercutting phase (but not reported).

Notice this specification implicitly constrains true asymmetric passthrough to be zero. The coefficient on $RACK$ in each of the equations is not allowed to vary depending on whether $RACK$ has increased or decreased from the previous period. This constraint is consistent with Eckert[2002], although his model is quite different than that here. Also for consistency, I do not include lagged rack price changes in the reported specification, but results are unchanged when I do.

The results are shown as specification (2) in Table 2. As predicted by the theory, the probability of continued undercutting decreases as $RACK_t$ increases or as $RETAIL_{t-1}$ decreases. Relenting becomes more likely. Conditional on relenting, the price increase is increasing in $RACK_t$ and decreasing in $RETAIL_{t-1}$. The (absolute value of the) expected price change in an undercutting phase also increases a modest amount when $RETAIL_{t-1}$ is higher or $RACK_t$ is lower. Interestingly, the two coefficients are always opposite in sign and, except with respect to the γ^U , of an insignificantly different magnitude. This shows that what matters for price decisions is a firm's position of the cycle, $RETAIL_{t-1} - RACK_t$, rather than either component individually.

To show how shocks are transmitted into overall expected retail prices (i.e. not conditional on regime) in the preliminary specification, I follow Eckert[2002] by calculating the partial derivatives of the unconditional price change with respect to the current rack price. The unconditional price change is simply given by

$$E(\Delta RETAIL_t) = prob(i_t = R | W_t^R, i_{t-1}) * E(\Delta RETAIL_t^R | X_t^R) \quad (5)$$

$$+ prob(i_t = U | W_t^U, i_{t-1}) * E(\Delta RETAIL_t^U | X_t^U) \quad (6)$$

which is then calculated for a series of nearby pairs of rack prices and differenced. Notice that this expectation correctly includes both realized and unrealized regimes – i.e. it considers the correct counterfactual – but *true* passthrough is still constrained to be identical across increases and decreases. Also, it is a static calculation in the sense that it involves only the immediate response at time t . Since the derivatives in general depend on values of $RETAIL_{t-1}$ and $RACK_t$, in Table 4, I report the partial derivatives for several values of $RACK_t$, fixing $RETAIL_{t-1}$ at its mean.

The results of Table 4 shows the derivative $\partial E(\Delta RETAIL_t)/\partial RACK_t$ is increasing and convex in the rack price, meaning that retail price is (increasingly) more sensitive to rack prices when $RACK$ is relatively higher (and therefore margin relatively lower). The main cause is regime shifting. The probability of switching to a relenting phase increases when $RACK$ increases, and should that switch occur, the expected price change will be much larger than had undercutting continued. Conversely, when $RACK$ decreases, the probability of switching to a relenting phase decreases and should that switch *not* occur, the expected price change will be much lower than had the firm switched to relenting.¹¹ The zero *true* asymmetric passthrough constraint forces these two effects to be equal in magnitude. The convexity occurs because these switching probabilities are most likely to be affected when margins are low.

Apparent asymmetric passthrough occurs because these two effects are differentially visible to real consumers. The large price increases triggered by a rack price increase are highly visible to the public. The large price increases that are *preempted* by negative rack price shocks are not visible to the public. The above results agree with Eckert[2002], though the methodology and sample cities are different. Interestingly, the result that retail prices are most sensitive when margins are small is also the same as Lewis[2004] for the U.S., who uses it as evidence for an entirely different model than Edgeworth Cycles or those suggested by

¹¹There is relatively less influence of $RACK$ on expected price changes, *conditional* on regime.

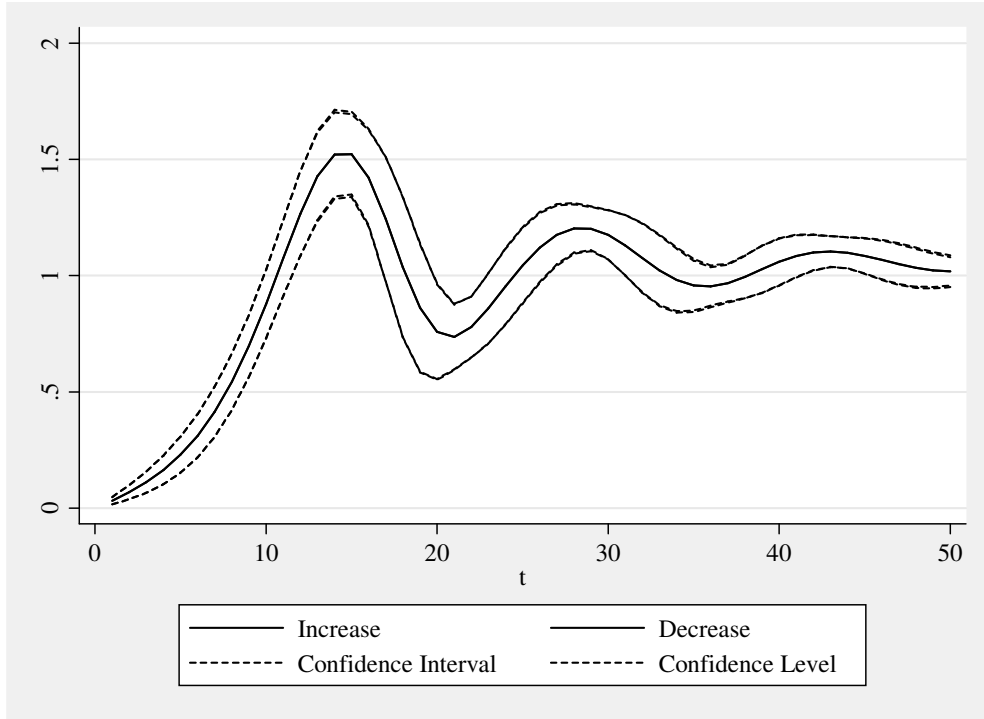


Figure 7: Cumulative Response Functions – Symmetric True Passthrough

Borenstein et al. I test the Lewis[2004] search model with respect to this particular market later.

The cumulative response functions for the symmetric case is reported in Figure 7. Examining the features of the rack-retail transmission will prove useful later on for interpreting results. These are the cumulative responses of retail prices to rack price shocks relative to *the correct counterfactual*.¹² Since true asymmetric passthrough is zero by constraint, the two cumulative response functions are necessarily identical. This occurs even though the Edgeworth Cycle itself self-generates asymmetric price movements.

First, the figure shows shocks are passed through quickly and overshoot complete passthrough after only eleven periods (5.5 days). This is because even small rack price changes can trigger

¹²I start $RETAIL_{t-1}$ at its mean but then allow it to evolve over time according to the estimated parameters of the cycle. Results are similar for any starting value of $RETAIL_{t-1}$, except that the initial rise in passthrough is delayed for higher values. This is consistent with the finding that $\partial E(\Delta RETAIL_{t-1})/\partial RACK_t$ increases with $RACK_t$ and decreases with $RETAIL_{t-1}$.

a change in the phase of the cycle (relative to the counterfactual). Relenting phase price changes average 5.57 cpl and undercuts average 0.43 cpl including zeros, so passthrough well over 100% is not surprising. Large price increases with little apparent justification at the wholesale level is a common source of suspicion in the press.

Second, the cumulative response functions oscillate over time – above and below 100% complete passthrough – before flattening out at 100% by the fiftieth period. This is because counterfactual prices – what would have happened to retail prices in the absence of a shock – is constantly changing in the background according to its cycle. Think of a large rack price increase that would with certainty cause a switch to a relenting phase. Actual price rises high above counterfactual. But a few periods later the actual price will have fallen some and the counterfactual price will leapfrog actual with its own cycle peak. Each time a new peak occurs in the one of these worlds, the price ordering reverses. Probablistically speaking, this causes an oscillation in expected prices. It eventually weakens as the variance on the distribution of possible states grow.

With these caveats in mind, I now turn to the primary analysis of true asymmetric passthrough under the correct counterfactual. If true asymmetric passthrough does not exist, the cumulative response functions will not significantly differ and we can conclude that all the apparent asymmetry is attributable to the Edgeworth Cycle. If it does exist, then we can further explore its possible causes.

I estimate a Markov switching regression model similar to that above except that I allow a full set of rack price lags into the model and allow the coefficients on those lags to differ depending on the direction of the rack price change. Specifically, I jointly estimate the system

$$\begin{aligned}\Delta RETAIL_{st} &= \beta_0^i + \beta_1^i RACK_{st} + \beta_2^i RETAIL_{st-1} \\ &\quad + \sum_{r=1}^{N/4-1} \beta_r^{i-} \Delta RACK_{s,t-4r}^- + \sum_{r=1}^{N/4-1} \beta_r^{i+} \Delta RACK_{s,t-4r}^+ + \beta_3^i IND + \varepsilon_{st}^i\end{aligned}$$

$$\lambda_{st}^{RR} = H(\theta_0^R)$$

$$\begin{aligned}\lambda_{st}^{UR} &= H(\theta_0^U + \theta_1^U RACK_{st} + \theta_2^U RETAIL_{st-1} \\ &\quad + \sum_{r=1}^{N/4-1} \theta_r^{U-} \Delta RACK_{st-4r}^- + \sum_{r=1}^{N/4-1} \theta_r^{U+} \Delta RACK_{st-4r}^+ + \theta_3^U IND)\end{aligned}$$

$$\gamma_{st}^R = H(\zeta_0^R)$$

$$\begin{aligned}\gamma_{st}^U &= H(\zeta_0^U + \zeta_1^U RACK_{st} + \zeta_2^U RETAIL_{st-1} \\ &\quad + \sum_{r=1}^{N/4-1} \zeta_r^{U-} \Delta RACK_{st-4r}^- + \sum_{r=1}^{N/4-1} \zeta_r^{U+} \Delta RACK_{st-4r}^+ + \zeta_3^U IND)\end{aligned}$$

where $i = R, U$, the number of lags is given by N and H is the logit function. IND is the firm type dummy. Recall λ^{RR} and γ^R are very close to zero and so only a constant is included in each. I include forty lags (twenty days) in the final specification.¹³ Because the price data is twelve hourly but rack prices are daily, consecutive (12-hour) rack price lags are highly collinear and I fit the coefficients $\Delta RACK^-$ and $\Delta RACK^+$ to a quadratic polynomial.¹⁴ For the same reason, I include every second two-period lagged difference of rack prices (i.e. daily rack) instead of every lagged difference.¹⁵

¹³Results do not change meaningfully with other reasonable lag lengths. Results do not meaningfully change when including sets of lagged retail price differences, which are generally insignificant.

¹⁴Results are similar when fit to a cubic, a fourth order polynomial or estimated freely (provided lags are sufficiently well separated to avoid multicollinearity and convergence problems.)

¹⁵Results are similar using every third three-period lag or every fourth four-period lag. Using longer gaps gradually dilutes the connection between rack price shocks and phase timing. Convergence problems can occur using every lagged difference.

Once the system coefficients are estimated, I consider three scenarios. In the first scenario (COUNTERFACTUAL), I fix current and all lagged values of the rack price $RACK_{t-r}$ at its mean in all periods. Retail prices, of course, continue to evolve with the cycle. This is the correct counterfactual against which we compare prices after a rack price shock. In the second scenario (RACK INCREASE), I impose a positive, permanent epsilon shock to $RACK$ beginning in some period q , and in the third scenario (RACK DECREASE), I impose a permanent negative epsilon shock at time q .

The expected price change at time t , $t > q$, is given by Equation 5 and calculated via a recursion algorithm

$$\begin{aligned}
E(\Delta RETAIL_t) = & [prob(i_{t-1} = R | W_{t-1}^R) * \lambda^{RR} + prob(i_{t-1} = U | W_{t-1}^U) * \lambda^{UR}] \quad (7) \\
& * E(\Delta RETAIL_t^R | X_t^R) \\
& + [prob(i_{t-1} = R | W_{t-1}^R) * \lambda^{RU} + prob(i_{t-1} = U | W_{t-1}^U) * \lambda^{UU}] \\
& * E(\Delta RETAIL_t^U | X_t^U)
\end{aligned}$$

noting that X_t^i and W_t^i contain $RETAIL_{s,t-1}$, $RACK_{s,t-r}$ and the history of lagged rack prices $\Delta RACK_{s,t-r}^+$ and $\Delta RACK_{s,t-r}^-$, and is thus appropriately path dependent.

After calculating unconditional expected price changes in every period under the three scenarios, I calculate the difference between the retail price path in the RACK INCREASE scenario and the COUNTERFACTUAL scenario for each period, and then between the COUNTERFACTUAL scenario and RACK DECREASE scenario for each period. The difference in differences yields the *true* asymmetric impact of rack price changes on retail price changes.

I perform five hundred simulations of each scenario in total. Each simulation is performed by taking a draw of the parameter vector from its estimated distribution and allowing prices




Figure 8: Cumulative Response Functions

to evolve along the cycle according to the parameter values for that simulation and given the inherent uncertainty contained in the σ 's. I take the mean of each outcome variable across simulations and calculate standard errors numerically.

Since the vectors of rack price coefficients are not readily interpretable, I focus again on the cumulative response functions of retail prices to cost shocks relative to the correct counterfactual. The response functions are reported in Table 5 and shown in Figures 8 through 11. $RETAIL_t$ is set to 45 for $t = 0$ (and evolves thereafter) and the shock occurs at the beginning of $t = 1$.¹⁶

Several recognizable features from the symmetric case re-emerge in Figure 8. First, passthrough rises quickly, whether it is an increase (mostly visible at the pump) or a decrease (not as visible at the pump), and overshoots 100% for both increases and decreases. It

¹⁶The results are similar for any starting values of $RETAIL$, with the first peak shifted right for higher values.




Figure 9: Cumulative Response Functions, $t = 1..10$

reaches 42% for increases and 69% for decreases by period 14. Second, the pattern of passthrough oscillates. We see that passthrough scales down to 87% and 65% for positive and negative shocks respectively by period 19 and then shortly thereafter exceeds 100% again. The oscillation straightens out over time and after the forty-fifth period (twenty three days and about three full cycles later) passthrough is not significantly different from 100% for either increases or decreases.

The most critical window for true asymmetric passthrough is the period up to the first peak in Figure 8. It is during this period that firms are first able to move prices upward following a positive shock or delay lowering them following a negative shock. Unlike in some crude or wholesale markets where inventory adjustment lags may play a role, any true asymmetric passthrough at the rack-retail is likely to first appear in this window. *Apparent* asymmetric passthrough should also be most noticeable early on.¹⁷

¹⁷Eckert[2002] finds with weekly data that the apparent asymmetric passthrough begins in the first week

Figure 9 shows a close-up of the periods up to the first peak. It shows evidence of *true* asymmetric passthrough. Up to the tenth period – the time in which the first “relent” is almost certain to occur – passthrough occurs significantly more quickly for positive shocks than for negative shocks.¹⁸ The difference reaches a maximum of 27 percentage points in the eighth period.

There are two sources for this result – differences in how rack price shocks affect the probabilities of regime switching, and differences in how rack price shocks affect expected price changes conditional on regime. The second effect turns out to be negligible, and the first effect drives the result. True asymmetric passthrough occurs because firms increase the probability of a relenting phase after a rack price increase more than the decrease the probability of a relenting phase after a rack price decrease. Said less formally, firms hurry relenting phases following a cost increase more than they delay relenting phases after a cost decrease. I discuss possible implications below.

One must be more cautious about interpretation beyond the first peak of Figure 8. The figure still shows the true passthrough rates at a point in time, of course, but path dependence now plays a role. The reason is that the first instance of asymmetric passthrough puts the cycles under the three scenarios out of phase, with cost shocks shifting cycle peaks forward or backward in time by different amounts. The difference in true passthrough naturally oscillates following the initial phase shift.¹⁹ This can be seen in Figure 10 which shows the difference in the cumulative response functions. True asymmetric passthrough reverberates between positive and negative values before becoming insignificant for good after the thirtieth period.

in Windsor. The duration of the cycle in Windsor is typically between two and four weeks.

¹⁸This result is robust across different specifications of the lag structure, as noted above.

¹⁹To see this most easily, reconsider figure 7 but imagine an extreme albeit unrealistic form of true asymmetric passthrough where firms pass through 100% of cost increases immediately. Assuming the actual-counterfactual spread always preserves this, the “increase” function would be represented in the figure as a horizontal line at one. If cost decreases are passed through according to the path depicted in the figure, then true asymmetric passthrough exists and oscillates – it is positive when the “decrease” function is above one and negative when below.

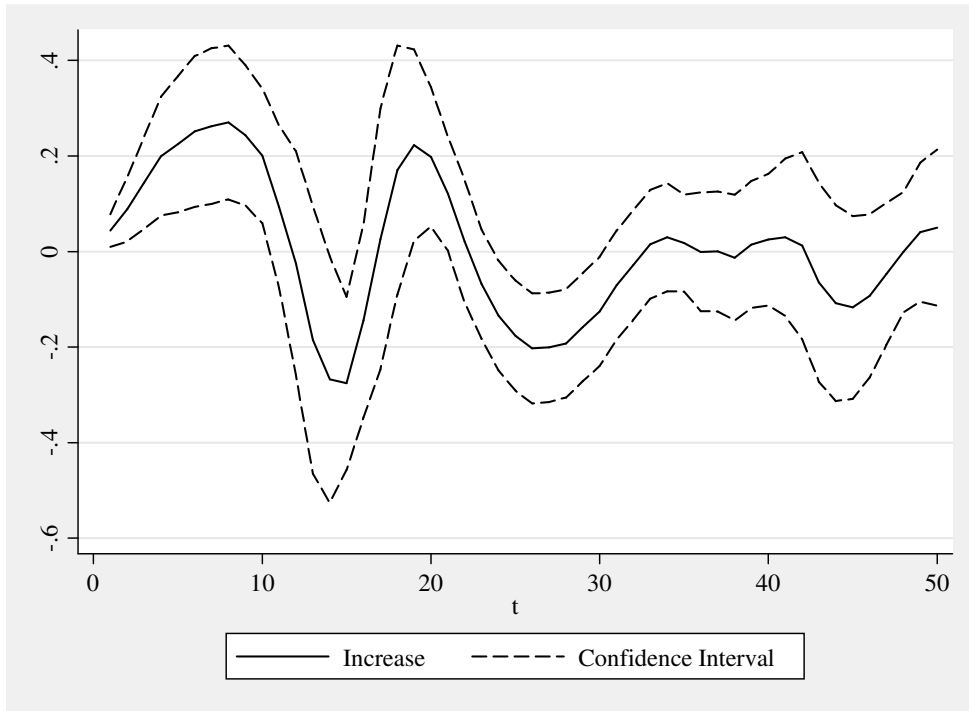


Figure 10: Difference in Cumulative Response Functions




Figure 11: Difference in Expenditures

The oscillation in true asymmetric passthrough implies that evaluating differences at a selected point in time, which is standard in the literature, is not the most appropriate for an overall welfare analysis. This is generally true whenever there are self-generating asymmetric price paths, like cycles. What matters is the entire history of prices over time. So rather than looking at cumulative price *changes*, I look at cumulative *prices*, or the difference in the integrals of the cumulative response functions. This would be equivalent to comparing *expenditures* after a positive and negative cost shock if one unit were purchased each period.²⁰

Figure 11 shows the difference in the expenditures following a positive versus a negative shock. Consumers pay significantly more early after the shock, becoming insignificant permanently after the twenty-fifth period (just under two full cycles). The sum of the additional payments that would need to be made after a one cent rack increase exceeds the extra savings after a one cent rack decrease – assuming one unit purchased per period – reaching a maximum of two cents in the tenth period. The long term difference in total expenditures eventually converges to a value of about 0.52 cents, getting very close to that value by the thirtieth period. I conclude there is a statistically significant degree of true asymmetric passthrough in the early periods under the correct counterfactual.

But is this difference economically significant? The answer seems to be no. One unit purchased at the average price of gasoline each period for 30 periods (15 days) gives a total expenditure of \$21.76, so the 0.53 cent difference is about 1/40th of 1% of expenditures over that short period. It represents 1/2 of 1% of net receipts (revenues over wholesale costs).²¹ These effects could still add up if there was a great deal of volatility in rack prices. However, the 15 day moving average of the sum of (the absolute value of) rack price shocks in the

²⁰In calling these “expenditures”, I put aside issues of intertemporal substitution which would work to benefit the price conscious consumer. I return to these issues later. Other reasonable weighting schemes for prices give similar results.

²¹I assume a discount factor of one, which should not pose a problem given the very short time frame involved (fifteen days). Using discount factors very close to one gives similar results.

sample was 4.24, of which about half are increases and half decreases. Therefore there is an average of 2.12 pairs of opposite-signed “one cent” shocks in 30 periods. The difference from all shocks is thus only 1/20th of 1% for expenditures, and 1% of net receipts. I conclude the true asymmetric passthrough is econometrically very close to zero. Contrast these numbers to the savings that consumers could easily achieve by following the three simple purchase timing rules I describe below: 3.2%, 6.5% and 7.7% reduction in expenditures respectively. The apparent asymmetric passthrough in the Toronto market is almost entirely due to the Edgeworth Cycles

7 Discussion

Even so, it is still worthwhile to explore some possible hypotheses for true asymmetric passthrough. Previous suggestions for (apparent) asymmetric passthrough in gasoline markets have included inventory distribution lags, focal pricing, covert collusion, and search. It has been difficult to pinpoint any single cause. The greatest single explanation for apparent asymmetric passthrough here is the cycles, and what follows is an exploration of true asymmetric passthrough only.

The first and common hypothesis is lags in inventory. However, distribution lags are not an issue at the rack to retail level and is easily ruled out.

A second hypothesis is focal pricing. Under focal pricing, firms hold prices at past levels following rack decreases, defecting only slowly over time. Meanwhile increases are passed through immediately. But for the standard focal price argument to hold, there needs to be some stability in prices for firms to focus on. This is far from true with the rapidly changing prices in the cycle. Even expected price changes conditional on regime and the probabilities of regime switching all respond quickly to increases and decreases. That is, even the second derivatives of prices are non-zero in every period.

A third hypothesis is the consumer search story of Lewis[2004]. In Lewis[2004], price

increases trigger search. Firms try to avoid this by raising prices only when necessitated by very low markups, and by raising them all the way back to the Nash level when they do. Search is not triggered when costs fall so prices fall only gradually. While a potential argument for apparent asymmetric passthrough (I return to this later) it is unlikely to be the cause of true asymmetric passthrough conditional on the cycle. True asymmetric passthrough is the difference in responses of prices after positive and negative shocks relative to the correct, but moving and unobserved, counterfactual. The story would require search to be triggered when the *change* in the expected price *change* (i.e. the second difference) relative to the counterfactual exceeds a certain positive threshold, which is unlikely. It should also discourage firms from raising prices when markups are not low. However, true asymmetric passthrough implies that firms *relent* to higher prices following a rack *decrease* relatively more quickly than under a no passthrough story (passthrough is less negative after a negative shock). In other words, firms do not just delay lower prices after a cost decrease – they actually raise them more quickly than under symmetric passthrough.²²

The folk theorem reminds us that there are infinitely many supergame strategies that could support a particular response pattern and so covert collusion can never be fully ruled out as a cause. However, this explanation would seem very unlikely given the complicated nature of the cycle. Successful collusion requires agreement, monitoring, and commitment to punish. The rapidly changing prices (about once a day per station) and competitive undercutting in this market would make complete pricing agreements on a daily basis heretically complex. Moreover, it is not the asymmetry inherent in the Edgeworth cycle that is of issue, but rather it is about asymmetric deviations in the response of the components of the cycle to wholesale shocks. Relative to the large price fluctuations from peak to trough,

²²Imagine the extreme case that firms ignore rack decreases and relent as they would in the counterfactual world, ie. zero true passthrough in the down direction, True asymmetric passthrough is thus strong – increases passed through much more quickly than decreases. Yet firms relent even though costs have fallen and markups are high, which triggers search. This is at odds with the Lewis model in which firms avoid triggering search unless very low markups make a price increase necessary.

the difference in profitability due to true asymmetric passthrough is very minor (<1% of net receipts). Intuitively, shocks causes a slight asymmetric wobble in the cycle itself. It is difficult to imagine firms would attempt to collude in this relatively minor and difficult dimension than by agreeing to stop the cycle entirely. In fact, the large increases that occur along the cycle without apparent wholesale justification is a major source of skepticism, and would best be avoided by colluding firms.

One last explanation for true asymmetric passthrough – which involves price leadership – remains quite plausible. Noel[2004] shows that coordinating price increases in the relenting phases of Edgeworth Cycles is a costly and risky proposition. There can be “delayed starts” and “false starts”. A delayed start occurs when other firms do not immediately follow the price increase by the first relenter. If the delay is extreme, the first relenter may even abandon its effort and return to the cycle bottom, creating a false start.²³ Even if other firms were to immediately match price increases, it is still costly to be first since the firm must still forgo its turn of being the low cost seller. This idea underlies the war of attrition. Therefore it is important to the first relenter that others follow as quickly as possible. A positive shock in rack prices may provide a natural impetus to help coordinate firms’ relenting behavior. If a negative shock to rack price does not unravel coordination to the same degree, true asymmetric passthrough can occur.

Hypotheses aside, the impact of true asymmetric passthrough in this market is small. The increase in consumer expenditures well under 1% is dwarfed by the savings consumers can achieve through the simple rules of thumb for timing purchases that I propose here. The fact is, with a little attention, the cycle peaks are very predictable. A simple and effective rule of thumb is to wait to buy gasoline until four days following the day prices have last risen.²⁴ This would save consumers 3.2% on gasoline expenditures immediately over purchasing at random times. A consumer can double the savings to 6.5% by watching for movements in the

²³Atkinson[2006] finds false starts are very common his detailed price data for Guelph, Ontario.

²⁴An increase of 5 cpl or more is a safe indicator.

rack price, which can be publicly attained. When the ex-tax price falls to within one cent of rack, she should buy.²⁵ A third technique does not require wholesale price information, and instead relies on slightly more careful observation of retail prices. It uses the fact it takes 24 - 36 hours after the first station has relented until the last one has. Therefore when the consumer sees a price difference of at least 5 cpl between two subset of stations, which in this market occurs about one day each week, it is a clear sign that a marketwide relenting phase is underway. By buying at one of lower priced stations, she can reduce her gasoline expenditures by 7.7% relative to randomly timed purchases.²⁶ The gains are substantial and readily attainable due to the predictability of the cycle. While these rules of thumb are finely tuned to this particular market, it would be easy to generate equivalent rules of thumb for other cycling markets as well.

Finally, and before concluding, it is interesting to contrast the models of (apparent) asymmetric passthrough studies by Eckert[2002] and Lewis[2004]. Both conclude that (apparent) asymmetric passthrough occurs because prices are relatively more sensitive to rack price shocks when markups are low rather than high. But each uses this finding to support a different model of retail competition at work. Lewis[2004] cites a model of costly consumer search for his U.S. study markets and Eckert[2002] cites Edgeworth Cycles for his study market of Windsor. This suggests a simple test to differentiate between these stories of apparent asymmetric passthrough for the market studied here. Under the Lewis search model, price increases should only occur once markups fall below a threshold level. Markups in turn fall because rack prices rise or because retail prices fall due to demand shocks. In the current paper, demand is slightly rising over the springtime sample period and cannot be the cause of low markups. Therefore, if the Lewis model is appropriate for this market, we should observe price increases only when markups are already low and a rack price increase occurs.

²⁵At the retail prices in effect over the sample, this corresponds to a difference between the posted retail price and the rack price of 30 cents.

²⁶If all consumers were to buy only at the trough, one can only speculate if the cycle would still occur.

Under the Edgeworth Cycle, large price increases can occur even when rack prices are flat or falling.

Therefore I test the null hypothesis that there is a zero probability of a relenting phase at a time t when rack prices have been constant in all periods $s < t$. To be conservative, I also check for relenting phases when rack prices have been falling steadily over time. I overwhelmingly reject the null hypothesis of no relenting phases when rack prices are constant (p value < 0.001), and reject it when rack prices are falling steadily by any reasonable amount (p value < 0.001 for each series of $\Delta RACK_t^- \in [-0.5, 0)$ for all t). That is, relenting phases occur provided rack prices fall less quickly than seven cents per week each week, which is well outside the range in the data. This supports Edgeworth Cycles over the consumer search model of Lewis as an explanation for much of the apparent asymmetric passthrough in this particular market.

8 Conclusion

The question of asymmetric passthrough is obscured by the presence of a strong asymmetric retail price cycle in some U.S. and Canadian gasoline markets, and in an increasingly large number of other gasoline markets worldwide. Conceptually, the asymmetry in the retail cycle is a separate phenomenon from any asymmetric response of retail prices to wholesale shocks. In this article, I test for true asymmetric passthrough – the asymmetry specifically due to increases vs. decreases in wholesale prices and net of the cycle – using a nested model of asymmetric passthrough and Edgeworth Cycles.

For the study market of Toronto, I find that a small degree of true asymmetric passthrough exists after removing the effect of the cycle. Firms hurry relenting phases after a rack price increase more than they delay them after a decrease. Passthrough quickly exceeds 100% of the shock for both increases and decreases before oscillating gradually toward 100%. True asymmetric passthrough occurs over the first five days with positive shocks being passed

through to retail prices more quickly than negative shocks, relative to the counterfactual. An expenditure measure is introduced to accommodate the oscillating pattern of passthrough over time. I find expenditures are significantly higher in the early periods but become insignificant thereafter. Collusion, inventory, and search explanations are unlikely, but efforts to overcome coordination problems in cycle resetting remains a plausible explanation. True asymmetric passthrough causes an increase of only 1/20th of one percent of expenditures. Compare this to simple rules of thumb under which consumers can save up to 7.7% of expenditures by timing purchases.

The general lesson of this article is that researchers going forward need to be careful of the choice of counterfactual in analyzing asymmetric passthrough. The discovery of Edgeworth Cycles in increasingly many gasoline markets around the world represents an increasingly important violation of the near universal assumption in the literature of a single steady state price. Many studies of asymmetric passthrough have been conducted for markets that are now known to have exhibited Edgeworth Cycles. The identifying assumption in these studies has been that tomorrow's price absent the shock would have been the same as today. The correct counterfactual, however, should be what tomorrow's price would have been in the absence of a shock today. In the large number of documented cycling markets, these are not the same thing.

9 References

- Atkinson, B. (2006) "Retail Gasoline Price Cycles: Evidence from Guelph, Ontario Using Bi-Hourly Station-Specific Retail Price Data", mimeo.
- Bachmeier, L and J.M. Griffin. (2003) "New Evidence on Asymmetric Gasoline Price Responses", *Review of Economics and Statistics* 85(3).
- Bacon, R. (1991). "Rockets and Feathers: the Asymmetric Speed of Adjustment of UK

Retail Gasoline Prices to Cost Changes” *Energy Economics* 13.

Balke, N. S. Brown and M. Yucel. “Crude Oil and Gasoline Prices: An Asymmetric Relationship?” *Federal Reserve Bank of Dallas Economic Review* 1st Qtr.

Borenstein, S., A. C. Cameron, R. Gilbert. (1997) “Do Gasoline Markets Respond Asymmetrically to Crude Oil Price Changes?” *Quarterly Journal of Economics* 112.

Castanias, R. and H. Johnson. (1993) “Gas Wars: Retail Gasoline Price Fluctuations”, *Review of Economics and Statistics* 75.

Cosslett, S. and L. Lee. (1985) “Serial Correlation in Latent Variable Models”, *Journal of Econometrics* 27.

Duffy-Deno, K. (1996) “Retail Price Asymmetries in Local Gasoline Markets” *Energy Economics* 18.

Eckert, A. (2002) “Retail Price Cycles and Response Asymmetry”, *Canadian Journal of Economics* 35, February.

Eckert, A. (2003) “Retail Price Cycles and Presence of Small Firms”, *International Journal of Industrial Organization* 21.

Ellison, G. (1994) “Theories of Cartel Stability and the Joint Executive Committee,” *RAND Journal of Economics* 25, Spring.

Godby, R., A. Lintner, T. Stengos and B. Wandschneider (2000) “Testing for Asymmetric Pricing in the Canadian Retail Gasoline Market”, *Energy Economics* 18.

Hendricks, K. (1996) Analysis and Opinion on Retail Gas Inquiry. Report prepared for Industry Canada.

Lee, L. and R. Porter. (1984) “Switching Regression Models with Imperfect Sample Separation Information – With an Application on Cartel Stability,” *Econometrica* 52.

Karrenbrock, J. (1991) “The Behavior of Retail Gasoline Prices: Symmetric or Not?” *Federal Reserve Bank of St. Louis Review* 73(4).

Kirchgassner, G. and K. Kubler (1992). “Symmetric or Asymmetric Price Adjustments in the Oil Market: An Empirical Analysis of the Relations between International and Domestic Prices in the Federal Republic of Germany, 1972-1989” *Energy Economics* 14.

Lerner, G. (1996) “Evaluation of the Six Resident’s Allegation of Price-Fixing in the Canadian Petroleum Industry,” Report prepared for the Competition Bureau, Industry Canada.

Lewis, M. (2004). “Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market”. Center for the Study of Energy Markets Working Paper 120.

Manning, D. (1991). “Petrol Prices, Oil Price Rises and Oil Price Falls: Some Evidence for the UK since 1972”, *Applied Economics* 23.

MJ Ervin & Associates. (1997) “Canadian Retail Petroleum Markets Study.” Report prepared for Industry Canada.

Maskin E., and J. Tirole. (1988) “A Theory of Dynamic Oligopoly II: Price Competition, Kinked Demand Curves and Edgeworth Cycles,” *Econometrica* 56.

Noel, M. (2007a) “Edgeworth Price Cycles, Cost-based Pricing and Sticky Pricing in Retail Gasoline Retail Markets,” *Review of Economics and Statistics*, 89:2, pp. 324-334.

Noel, M. (2007b) “Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market,” *Journal of Industrial Economics* 55:1, pp. 69-92.

Noel, M. (2004) “Edgeworth Price Cycles and Focal Prices: Computational Dynamic Markov Equilibria” UCSD Working Paper, 2004-12.

Radchenko, S. (2005). “Anticipated and Unanticipated Effects of Crude Oil Prices and

Gasoline Inventory Changes on Gasoline Prices.” UNC Charlotte mimeo.

Reilly, B. and R. Witt. (1998) “Petrol Price Asymmetries Revisited” *Energy Economics* 20.

Slade, M. (1987) “Interfirm Rivalry in a Repeated Game: An Empirical Test of Tacit Collusion,” *Journal of Industrial Economics* 35.

Slade, M. (1989) “Price Wars in Price-Setting Supergames,” *Economica* 56.

Slade, M. (1992) “Vancouver’s Gasoline Price Wars: An Empirical Exercise in Uncovering Supergame Strategies,” *Review of Economic Studies* 59.

Wang, Z. (2005). “Strategy, Timing and Oligopoly Pricing: Evidence from a Natural Experiment in Gasoline Markets”. Monash University mimeo.

Zhang, X. and J. Feng. (2005) “Price Cycles in Online Advertising Auctions”. MIT mimeo.

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

In Canadian cents per liter.

Table 2: Regime and Switching Probability Results

Specification:	(1)	(2)
RELENTING PHASE (dep. var. = $\Delta RETAIL_{st}$)		
$\alpha^R = E(\Delta RETAIL_{st} I_{st} = "R")$	5.576 (0.083)	
$RETAIL_{st-1}$		-0.887 (0.029)
$RACK_{st}$		0.820 (0.028)
σ_R	1.650 (0.066)	0.729 (0.057)
STICKY PRICING – RELENTING PHASE		
$\gamma^R = \Pr(J_{mt} = 1 I_{st} = "R")$	0.000 (0.000)	
UNDERCUTTING PHASE (dep. var. = $\Delta RETAIL_{st}$)		
$\alpha^U = E(\Delta RETAIL_{st} I_{st} = "U")$	-0.751 (0.008)	
$RETAIL_{st-1}$		-0.043 (0.004)
$RACK_{st}$		0.052 (0.005)
σ_U	0.459 (0.009)	0.419 (0.008)
STICKY PRICING – UNDERCUTTING PHASE		
$\gamma^U = \Pr(J_{mt} = 1 I_{st} = "U")$	0.429 (0.007)	
$RETAIL_{st-1}$		0.057 (0.015)
$RACK_{st-1}$		-0.019 (0.015)
SWITCHING PROBS. ($\lambda^{ik} = \Pr(I_{st} = k I_{s,t-1} = i)$)		
$\lambda^{RR}(R \rightarrow R)$	0.008 (0.004)	
$\lambda^{UR}(U \rightarrow R)$	0.078 (0.001)	
$RETAIL_{st-1}$ in λ^{RU}		-0.749 (0.036)
$RACK_{st}$ in λ^{RU}		0.777 (0.037)

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 delta method.

Table 3: Cycle Characteristics

Specification	(1)
Relenting Phase Duration	1.008 (0.004)
Undercutting Phase Duration	12.780 (0.291)
Cycle period	13.788 (0.291)
Asymmetry	12.680 (0.291)
Cycle Amplitude	5.619 (0.082)

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 delta method.

Table 4: Partial Derivatives of Retail Price Changes to Rack Prices

$RACK_{st}$	$\partial E(\Delta RETAIL_t) / \partial RACK_t$
37.6	0.045 (0.004)
38.2	0.061 (0.004)
38.8	0.087 (0.005)
39.4	0.130 (0.007)
40.0	0.197 (0.010)
40.6	0.295 (0.016)
41.2	0.427 (0.025)
41.8	0.580 (0.034)
42.4	0.729 (0.039)

Table 5: Cumulative Response Functions

	Cumul. Diff. + vs NC	Cumul. Diff. NC vs -	Asym.	S.E.		Cumul. Diff. + vs NC	Cumul. Diff. NC vs -	Asym.	S.E.
1	0.057	0.012	0.044	(0.017)	26	1.198	1.401	-0.202	(0.058)
2	0.114	0.025	0.088	(0.034)	27	1.246	1.446	-0.200	(0.058)
3	0.190	0.046	0.143	(0.049)	28	1.258	1.451	-0.192	(0.057)
4	0.268	0.069	0.199	(0.063)	29	1.258	1.417	-0.158	(0.057)
5	0.329	0.104	0.224	(0.072)	30	1.224	1.349	-0.125	(0.058)
6	0.399	0.147	0.251	(0.080)	31	1.188	1.259	-0.070	(0.058)
7	0.481	0.219	0.262	(0.083)	32	1.143	1.171	-0.028	(0.058)
8	0.583	0.313	0.270	(0.082)	33	1.120	1.105	0.015	(0.058)
9	0.711	0.467	0.243	(0.074)	34	1.111	1.081	0.029	(0.057)
10	0.863	0.663	0.201	(0.071)	35	1.136	1.118	0.017	(0.051)
11	1.043	0.950	0.093	(0.086)	36	1.173	1.174	-0.001	(0.063)
12	1.220	1.245	-0.024	(0.119)	37	1.234	1.234	-0.001	(0.064)
13	1.368	1.553	-0.184	(0.142)	38	1.288	1.301	-0.12	(0.067)
14	1.422	1.690	-0.266	(0.131)	39	1.352	1.338	0.014	(0.067)
15	1.371	1.647	-0.274	(0.092)	40	1.392	1.367	0.024	(0.070)
16	1.218	1.363	-0.144	(0.103)	41	1.350	1.320	0.029	(0.083)
17	1.049	1.022	0.026	(0.139)	42	1.283	1.267	0.012	(0.099)
18	0.921	0.750	0.170	(0.133)	43	1.145	1.209	-0.064	(0.106)
19	0.868	0.646	0.222	(0.102)	44	1.046	1.148	-0.108	(0.104)
20	0.870	0.672	0.198	(0.074)	45	0.971	1.087	-0.117	(0.097)
21	0.906	0.784	0.121	(0.060)	46	0.939	1.031	-0.092	(0.087)
22	0.953	0.932	-0.020	(0.065)	47	0.939	0.986	-0.048	(0.075)
23	1.016	1.085	-0.068	(0.058)	48	0.953	0.954	-0.001	(0.064)
24	1.078	1.211	-0.133	(0.058)	49	0.977	0.937	0.040	(0.074)
25	1.147	1.323	-0.176	(0.058)	50	1.002	0.950	0.050	(0.083)