A more recent version of this article has been accepted for publication in:
International Journal of Industrial Organization

Do Edgeworth Price Cycles Lead to Higher or Lower Prices?

Michael D. Noel*

August 4, 2015

Abstract

A recent literature seeks to understand the causes of the high-frequency, asymmetric retail price cycles observed in many retail gasoline markets. However, much less attention has been given to the *effects* of the cycles, in particular, whether the cycles lead to higher or lower prices and margins. The leading theory for the underlying cause of the price cycles, Edgeworth price cycles, is silent on the issue. The challenge in addressing this most important question has been the difficulty in isolating cycle effects from other confounding factors, especially market structure. In this article, I exploit a unique natural experiment to isolate the effect of cycles - a refinery fire that, in a matter of days, halted cycles that had previously persisted for decades. I find that Edgeworth price cycles lead to lower prices and lower margins. I conclude with implications for competition policy.

1 Introduction

In retail gasoline markets, where firms sell a relatively homogeneous good, one might expect prices to closely follow costs not only in the long run but also in the short run. One stark exception to this occurs in markets where gasoline prices follow a high-frequency and asymmetric "sawtooth-like" cycle. In a price cycle, prices rise very quickly, often by five or ten percent, in a few days or even a few hours. Prices then rachet down for the next week or two or even perhaps just the rest of the day, before increasing again. The cycle repeats over and over even in the absence of any changes in wholesale prices.

<Figure 1 about here>

Figure 1 shows an example of a weekly cycle and a daily cycle in the cities of Windsor, Canada and Toronto, Canada, respectively. The former uses daily data and the latter uses data collected

*Michael D. Noel, Associate Professor, Department of Economics, Texas Tech University, Lubbock, Texas, michael.noel@ttu.edu.

four times a day. Cycles broadly similar to these have been observed in markets across Canada, the Midwest U.S., Australia, and numerous European countries. They were also present in southern and western U.S. cities in the past.

A growing literature has emerged to try to understand the causes of these cycles. The standard theory is that of Edgeworth price cycles, the outcome of a dynamic oligopoly game in which firms using Markov strategies sell homogeneous goods and compete in prices (Maskin and Tirole (1988)). Various authors have empirically tested the predictions of the theory and of extensions of it by Eckert (2003) and Noel (2008), and generally find support for the Edgeworth price cycle theory (Eckert (2003), Noel (2007a, 2007b), Atkinson (2009), Wang (2009a, 2009b), Doyle, Muehlegger and Samphantharak (2010), Lewis (2012), and others). However, the theory is silent on the most important and policy relevant question of all – do cycles lead to higher or lower prices overall?

Recently, a new strand of literature has suggested that the cycles may in fact not be competitive at all but rather collusive in nature (Foros and Steen (2013), de Roos and Katayama (2013), Byrne and de Roos (2014)). Competition authorities around the world have shared this concern and have investigated firms in cities that experience cycles in search of explicit collusion (e.g. U.S., Canada, Germany, Norway, and Australia). If cycles were in fact the result of collusion, consistent with this concern, then they should surely lead to higher prices at a minimum.

So do price cycles actually lead to higher prices or lower prices? It is ultimately an empirical question, but surprisingly, in spite of its centrality to the literature, it has never been convincingly answered. The reason is because of difficulty with identification of a causal effect. Cycles typically persist for decades, and authors must often rely on cross-sectional comparisons across cities with and without, subject to all the usual and well-known omitted variables biases (Noel (2002), Doyle et al. (2010)). For example, it has been shown that cycles are more likely to appear in markets with more price aggressive independent firms (Eckert (2003), Noel (2007a, 2007b), Doyle et al. (2010), Lewis (2012)). But since these firms have a direct effect on price levels separate from any effects of the cycles, effects of market structure and of the presence of cycles are difficult to separate in a cross-section. Authors note this and caution against a causal interpretation of cross-sectional comparisons. Zimmerman, Yun, and Taylor (2013) were first to use a panel dataset to examine cycle effects on prices, but as the growth of several independent chains in the Midwest coincide with the appearance of cycles, the authors again caution against a causal interpretation.

In this article, I exploit a unique natural experiment that, combined with a difference-in-differences panel framework, enables me to estimate a plausibly causal effect of Edgeworth price cycles on price and margin levels for the first time. The natural experiment surrounds a refinery fire that occurred on February 15, 2007 in Nanticoke, Ontario, Canada, on the shore of Lake Erie. The fire caused a temporary gasoline supply shock and, within days, retail price cycles that had persisted for decades in several southern Ontario cities came to an abrupt halt. When the refinery and surrounding markets returned to normal operation a month later, the only thing different was that the decades-old retail price cycles in those cities were no more. In contrast, other cities further away from the refinery that had cycles before February 2007 continued to have cycles after. Cities that did not have cycles before February 2007 continued to not have them after.

This yields clear treatment and control groups and a clear break in the equilibrium type that lends itself well to a difference-in-differences framework and estimation of causal effects. The break in equilibrium type means I can compare prices and margins in the same cities, first with and then without cycles, holding fixed city-specific unobservables including market structure, all within months of each other. The difference-in-differences panel framework means I can control for changes and trends in prices and margins due to unobservable common shocks, which turns out to be important since prices were on the rise generally right around the time of the fire. The resulting estimates lend themselves well to a causal interpretation, avoiding the usual concerns in the literature about omitted variable bias from a straight cross-sectional comparison or a simple before-and-after case study.

The sudden loss of cycles in several Ontario cities was first described by Atkinson, Eckert and West (2014). They examine its possible causes (though not its effects) and conclude the fire was clearly responsible for the change – the collapse of cycles was immediate. The initial disappearance of the cycles and the fact they did not re-emerge when supply conditions normalized is consistent with the multiple equilibria nature of the Edgeworth price cycles theory in the face of temporary capacity constraints (Noel (2008)). Thus I observe prices and margins under both cycling and non-cycling equilibria in the absence of other confounding factors.

Getting a measure of the causal effects of cycles is critical for improving our understanding of cycling markets and for competition policy. First, it answers one of the most important but elusive questions in the literature. Are cycles good or bad for competition? Efficiency and welfare analysis ultimately care about whether prices are competitively set. Second, if I were to find that cycles actually lowered prices, it would challenge the predicate often used in investigations that cycles are inherently undesirable. Investigations tend to focus on searching for evidence of explicit collusion underlying the rapid sequence of price increases that would violate existing antitrust laws. Understandably, price increases viewed in isolation can appear unjustified absent any contemporaneous increases in costs. However, less attention is afforded the undercutting process that dominates the rest of the cycle and, importantly, to evaluating the entire cycle process as a whole. This article highlights the importance of evaluating the cycle as a whole.

To preview results, I find that the presence of Edgeworth price cycles in the sample markets before the fire *lowered* prices and margins. In other words, the disappearance of those cycles, all else equal, led to higher prices and margins. The results are not consistent with a collusive story for the existence of cycles. It is, however, consistent with a beneficial view of Edgeworth price cycles and suggests that regulatory efforts to eliminate the cycles may in fact do more harm than good. The results are robust to a wide range of alternative specifications and assumptions.

I proceed as follows. Section 2 reviews the relevant literature, Section 3 discusses the data and methodology, Section 4 reports results, and Section 5 contains further discussion. Section 6 concludes.

2 Literature & Background

The concept of an Edgeworth price cycle dates back to Edgeworth (1925) and was formalized by Maskin & Tirole (1988) as the outcome of a dynamic oligopoly game with homogeneous goods, alternating moves, and Markov strategies. Markov strategies depend only on the payoff relevant state which, in this context, includes the current price of the other firm, current cost, and any other state variables such as demand.

Let $V^1(p^2_{t-1})$ be the Firm 1's value function when Firm 2 adjusted its price to p^2_{t-1} in the previous period, and Firm 1 adjusts its price in the current period. Let $W^1(p^1_{s-1})$ be Firm 1's value function when it has set price p^1_{s-1} in the previous period and Firm 2 is about to adjust its price. V^1 and W^1 are written:

$$V^1(p^2_{t-1}) = \max_{p_t} \left[\pi^1_t(p_t, p^2_{t-1}, c) + \delta W^1(p_t)\right] \tag{1}$$

$$W^1(p^1_{s-1}) = \mathop{\mathrm{E}}_{p_s} \left[\pi^1_s(p^1_{s-1}, p_s, c) + \delta V^1(p_s)\right] \tag{2}$$

where π^i_t is the t-period profit of firm i and δ is the common discount factor. The expectation W^1 is taken with respect to the distribution of R^2, the best response function for Firm 2. Similar equations are found for V^2 and W^2.

<Figure 2 about here>

Maskin and Tirole show that two sets of Markov Perfect Equilibria (MPE) are possible even under the same demand conditions – "focal price" equilibria and "Edgeworth price cycle" equilibria. Focal price equilibria are characterized by constant prices or margins over time. Edgeworth price cycle equilibria yield best response functions of the form:

$$R^i(p^j) = \begin{cases} \overline{p} & \text{for } p^j > \overline{p} \\ p^j - k & \text{for } \overline{p} \geqslant p^j > \underline{p} \\ c & \text{for } \underline{p} \geqslant p^j > c \\ c & \text{with probability } \mu(\delta) \text{ for } p^j = c \\ \overline{p} + k & \text{with probability } 1 - \mu(\delta) \text{ for } p^j = c \\ c & p^j < c \end{cases} \tag{3}$$

where k represents one step on the discrete price grid and c is the constant marginal cost. The equilibrium price path exhibits an asymmetric sawtooth-shaped price cycle. Figure 2 shows an example of an Edgeworth price cycle with $c = 0$ (Noel and Chu (2015)).

The practical mechanics are straightforward. Firms undercut one another until the price gets relatively close to marginal cost, at which point one firm drops its price immediately to marginal cost. Firms then play a war of attrition, each playing a mixing strategy, mixing between raising the price back to $\overline{p} + k$ and maintaining price equal to marginal cost. By holding at marginal cost, the firm creates the possibility that the other firm will raise price first, but runs the risk that the other firm will hold at cost too, meaning the first firm will find itself in the same zero profit situation two periods hence. Eventually, one firm does raise price to $\overline{p} + k$, the other undercuts, and tit-for-tat undercutting begins again. In asymmetric versions of the model, a single firm can emerge as the sole price leader in increasing prices (Noel (2008)).

Allvine and Patterson (1974) first noted asymmetric retail gasoline price cycles in southern and western U.S. cities in the late 1960s and early 1970s, and Castanias & Johnson (1993) first noted the similarity between the Los Angeles cycles and the then new model of Edgeworth price cycles.[1]

Beyond the price asymmetry, the basic theory – which assumes a homogeneous symmetric duopoly with constant costs, offered little in the way of testable predictions. With the rediscovery of asymmetric cycles in Canadian markets from at least the 1980s, several authors extended the theory with the realities of retail gasoline markets in mind. Eckert (2003) extends the Bertrand setting to allow for an uneven split of the market even when prices were equal – with the "larger" firm getting the larger share at equal prices. He predicts cycles are more likely with more asymmetric firms, with the larger firm initiating price increases and the smaller one initiating price decreases. Noel (2008) extends the model in a variety of other dimensions, including stochastic marginal costs, capacity constrained models, models of product differentiation, triopoly, and other kinds of asymmetric equilibria. The author finds cycles are robust to most settings (but not if products are too differentiated or capacity constraints too tight) and finds that the shapes of cycles depend on model parameters in testable ways. Noel also predicts that "false starts" – when firms abandon price increases after others do not follow quickly enough – are potentially part of the triopoly equilibrium process.

Empirical papers largely support the Edgeworth price cycle theory. Eckert (2003) finds that price volatility was greater in Canadian cities that had more independent firm presence, and Noel (2007a) finds that asymmetric price cycles in particular were more likely. Noel (2007b) and Atkinson (2009) use high frequency datasets to show that large firms initiated price increases and smaller

[1] Edgeworth price cycles have rarely been seen in pricing patterns outside retail gasoline. For exceptions (internet keyword auction markets), see Zhang (2006) and Edelman and Ostrovsky (2007).

independents initiated undercutting for the cities of Toronto and Guelph, Ontario, respectively. Byrne, Leslie & Ware (2015) find a similar result across a wider cross section of small Ontario towns. In his bi-hourly data, Atkinson also observes the false starts predicted by Noel (2008).

After a long absence, cycles reappeared in the U.S. in many Midwest cities in 2000 (Zimmerman et al. (2013)). Doyle et al. (2010) find they were more likely in cities with more large independent firms and with more attached ancillary businesses such as convenience stores. Lewis (2012) finds that large firms initiate price increases but that the large firms in this case were two large independent chains – Speedway and Quik Trip.

All major cities in Australia exhibit price cycles and Wang (2009a, 2009b) documents the cycles in Perth. Wang (2009a) shows firm-level cross price elasticities to be very high, consistent with the Edgeworth price cycles model, and Wang (2009b) finds that major firms BP, Caltex, and Mobil played mixed strategies, also as predicted in the symmetric model. He exploits the simultaneity feature of the price regulation in effect in Perth known as FuelWatch – which requires firms to "simultaneously" choose and commit to prices a day in advance – to estimate mixing probabilities. Bloch & Wills-Johnson (2010) show that the rate of undercutting in Perth is faster in neighborhoods with more competition, also consistent with the model.

Some authors express concern that the cycles are collusive in nature or contain collusive elements.[2] Studies typically focus on documenting the speed and uniformity of price increases in the relenting phase of the cycle. Foros and Steen (2013) argue that the cycles in Norway are hardly consistent with competition as evidenced by the fact that price increases in Norway occur within hours of each other on the same day of the week.[3] Byrne and de Roos (2014) examine Perth and document that price restoration phases there became smoother over time, faster and with more homogeneous actions by stations.[4] As noted by Noel (2008) in his theoretical extension, the degree of difficulty in leading restoration phases and having others follow can vary and lead to extended

[2] Byrne, Leslie, and Ware (2015), which examine cycles in Ontario, state "we stop short of attempting to detect collusion in cycling gasoline markets empirically. Such an ambitious task...is clearly an important area for future research." Interestingly, the results of this article suggest that potential widespread collusion that would be responsible for the creation of the cycles themselves is unlikely to exist.

[3] Similar day-of-the-week price increases occurred in Australian cities (Noel and Chu (2015)). The feature was popular with Australian consumers since it allowed them to more easily predict the troughs.

[4] De Roos & Katayama (2013) argue the Maskin and Tirole alternating moves framework can be trivially ruled out in this city since stations do not decrease price sequentially one station per day as in the model.

wars of attrition and disruptions.

While the literature has focused on the causes of the cycles, much less attention has been given to directly measuring their effects. Obviously, if cycles were collusive, the effects would be negative. If cycles are simply Edgeworth price cycles, the theory is ambiguous on effects.

As cycles tend to persist for decades, often before the start of available data, typically only cross-sectional comparisons of prices and margins across cities have been possible. Noel (2002) reports prices and margins were 2.5 Canadian cents per liter (cpl) lower in Canadian cities with cycles than in cities without. Doyle et al. (2010) find prices are 1.5 to 3 U.S. cents per gallon lower in the U.S. cities with cycles than in those without (or 1-2% at a then average price of $1.52). Both studies caution against causal interpretations. In one before-and-after comparison, Wang (2009b) finds prices in Perth fell by 0.9-1.8 Australian cents per liter when the cycles temporarily faltered for several months immediately following the introduction of the FuelWatch price pre-commitment and pre-notification regulation. However, it is not clear if the lower prices during this disequilibrium period are representative of what a long term non-cyclical equilibrium would have looked like. In another before-and-after, Atkinson, Eckert, and West (2014) note that margins increased in the city of Toronto after the Nanticoke fire, but since margins were generally increasing in other cities, and since Toronto's almost perfectly constant margin pattern post-fire was potentially unusual (as discussed in detail below), a causal interpretation was elusive.

The closest relative to this study is that of Zimmerman, Yun and Taylor (2013) which examine a panel of prices before and after retail price cycles began in the Midwest U.S. in 2000. Comparing price changes in Midwest cities after cycles first began to changes in cities that never had cycles, they find that prices in cities with cycles decreased by 0.75 to 1 U.S. cent per gallon. The one concern is that the panel is a long panel and the start of the Midwest cycles in 2000 roughly coincides with the expansion of large independent chains there. Lewis (2012) finds that two such independent chains, Speedway and Quik Trip, are largely responsible for leading price increases each week and generating the cycle. It cannot be ruled out that the estimated price effects combine both the effects of cycles and also the direct effects from changes in market structure. The authors caution against a causal interpretation.

Noel (2008) takes a computational approach to theoretically compare prices under various cycling and non-cycling MPE. The author finds Edgeworth price cycles routinely produce lower prices than stable pricing in the computed equilibria. While suggestive, the set of equilibria is not comprehensive and the results cannot be generalized to all possible MPE.

Several authors have also addressed the potential price effects of cycles indirectly. Lewis (2009) finds that Edgeworth price cycles sped up market recovery following Hurricane Rita and Lewis and Noel (2011) find that cycles speed up market recovery after general cost shocks by a factor of two to three times, by reducing the "rockets and feathers" effect.[5] Noel (2012) and Noel and Chu (2015) show, for Canada and Australia respectively, that cycle troughs can be predicted, allowing interested consumers to systematically purchase at prices 4% below the unweighted average price given in market data. Studies supporting the collusive view tend to focus on describing patterns in the speed and sequence of restoration phases and inferences about price effects from these studies are by their nature indirect.

The cessation of cycles after the Nanticoke refinery fire is consistent with the theoretical model of Edgeworth price cycles. Noel (2008) shows that relatively strong capacity constraints reduce the gain to undercutting, destabilize cycles, and lead to more stable, non-cyclical price setting. The absence of a return back to price cycles after the capacity constraint problems were resolved is also consistent with the multiple equilibria nature of the model. In terms of a collusive story, it is less clear why a negative supply shock should make collusion less likely, but the multiple equilibria caveat still applies.

3 Data and Methodology

The Nanticoke refinery is one of several refineries serving southern Ontario cities and to a much lesser extent, cities in then neighboring province of Quebec. It does not serve Western Canada.

To estimate the causal effects of cycles on prices and margins, I use daily market-average data on retail prices and wholesale rack prices for a cross-section of nine major Canadian cities over a period of two years. The cities include four from the province of Ontario (Toronto, Ottawa, London, Windsor), two from Quebec (Montreal, Quebec City) and three from Western Canada (Vancouver,

[5] Eckert (2002) and Noel (2009) show how Edgeworth price cycles and the so-called rockets and feathers effect (e.g. Borenstein, Cameron, and Gilbert (1997), and many others) are interdependent.

Calgary, Winnipeg).[6]. The data was purchased from GasBuddy.com and are derived from all price reports submitted to the company's website by a voluntary network of members or "spotters". The number of daily price reports range from a few dozen per day in Windsor and Quebec City to over a thousand per day in Toronto. The average is 237 reports per day. The sample spans from February 15, 2006 to February 14, 2008, one year on either side of the February 15th, 2007 fire.[7] I supplement this data with four-times-daily city-average price data, as discussed below, also from GasBuddy. Intraday periods are 6am-10am, 10am-2pm, 2pm-6pm, and 6-12am. The percentage of price reports in each intraday period are 35%, 19%, 18%, and 28% respectively.

Atkinson (2008) finds the GasBuddy data to be largely reliable (comparing the price data to self-collected prices from Guelph stations), and it is widely used by researchers.[8] A nice feature of the GasBuddy data is that stations that are observed and visited more frequently are also sampled more frequently, essentially weighting prices by a proxy for quantity or popularity. Price reports as a proxy for quantity have recently been advanced by Lewis and Marvel (2011) and Byrne, Leslie, and Ware (2015).

I match the daily data with average daily posted rack prices (wholesale prices), for each city, collected from refiners by Kent Marketing Ltd.[9] The rack price is the best available measure of

[6] These are nine of the largest sixteen cities in Canada by population (2006 figures), including eight of the top ten. Of the other top-sixteen cities, western cities of Edmonton and Victoria were excluded because of proximity to other control cities, Halifax because gasoline was price capped, and the remaining ones because of proximity to Toronto in the Golden Horseshoe area. I follow the standard practice of focusing on major centers, given that data for smaller towns are more thin, especially near the start of the sample in early 2006 (prior to the proliferation of smartphones). One study that uses GasBuddy data for Ontario since the smartphone era and includes smaller towns as well as larger ones is Byrne and Leslie (2013).

[7] Data for Montreal and Quebec City are available from June 2006 only.

[8] Possible reporting errors in the GasBuddy data are unlikely to be of great concern, especially given Atkinson (2008) and the difference-in-differences framework used here. If spotters just make random errors (in the dependent variable) it adds noise but not bias. If instead spotters systematically over- or under-estimated prices everywhere, the bias would simply be swept into the constant term and differenced out of the analysis. If the errors were systematically different across treatment and control groups, or across before and after periods, they would again be harmlessly swept into the city dummies or before/after period dummies. To affect results, it would have to be the case that the bias occurs in only one of the two treatment groups for only one of the two time periods (or in all except for one group and one period), a strong requirement. Even in such a case, the degree of misreporting would need to be extreme to matter. To negate the effects I find below, at least 10% of all spotters would have to systematically under-report or systematically over-report prices by a full 10 to 15 cpl (when prices are a dollar) in every price report for every station within that one group during just that one period, while reverting to unbiased reporting otherwise. It is not realistic. A second concern may be that some reports, while not incorrect, are late, e.g. a price observed in the evening is not reported until the next morning. However, this also cannot realistically affect results. Even if every single spotter reported every single price a full 24 hours late, the treatment effects presented here change by only 1/1000th of a penny.

[9] No rack price is published for Windsor, so I use that of the nearest city.

wholesale price. Also, as sample markets are all proximate to available U.S. supplies, the rack price can reasonably be modelled as exogenous. Prices are adjusted by federal and provincial taxes with information reported by Kent Marketing. Margins are calculated as the difference between retail and rack price net of taxes. Summary statistics for the daily dataset are given in Table 1.

To implement a difference-in-differences model, I must first divide the markets into treatment markets and control markets.[10] Treatment markets are those whose cycles ceased immediately after the February 15, 2007 fire. Control markets are those whose previous equilibrium type, either cycling or non-cycling, remained the same. Although the division of markets into treatment and control groups is obvious by observation of the sample data, I provide some statistical metrics for reference as well.

I first examine daily data for evidence of weekly price cycles and any fire-induced stoppage in those weekly cycles. Weekly cycles are evident in two cities. Figure 1 clearly shows the weekly price cycles for the city of Windsor. The graph for Montreal (not shown) is similarly clear.

Lewis (2009) proposes a median price change metric for categorizing cities with and without cycles. With data frequent enough to observe the asymmetry, markets with cycles should exhibit a more negative median price change than cities without, since there will be relatively few large price increases but many small price decreases along cycles.[11] I find the median price change was -0.74 cpl in Montreal and -0.85 cpl in Windsor before the February 15th fire. The median price change after the fire and after the supply shock had passed (using a date of March 26th, see below), was -0.66 cpl in Montreal and -0.45 cpl in Windsor. While the intensity of the Windsor cycles lessened to an extent after the fire, a point to which I return, asymmetric cycles are clearly evident both before and after the fire. In all other cities, before and after, the median price change in the daily data is 0.1 cpl or less in absolute value.

It is well known that prior to the fire, Vancouver and numerous cities in Ontario experienced daily rather than weekly or longer period cycles. Daily cycle asymmetry can be difficult to observe

[10] It is reasonable to treat each city discretely as either cycling or not cycling. Even in larger cities such as Toronto or Vancouver, a single cyclical pattern tends to flourish across the entire city, with peaks and troughs lining up in time. This is largely due to firms coordinating price increases across its own stations citywide (Noel (2007b) for Toronto, Eckert and West (2004) for Vancouver, Lewis (2012) for the U.S.)

[11] The threshold for the median price change must still be determined by the observer. It will depend on the frequency of data and whether prices are trending higher or lower.

in daily data and will not show a large negative median price change. Therefore, to verify the existence of daily cycles in Vancouver and Ontario cities prior to February 15, 2007, and the subsequent cessation of daily cycles in Ontario cities and continuation of cycles in Vancouver, I obtain price data at a four-times-daily frequency for Vancouver and the Ontario cities. Once equilibrium changes in daily cycles are identified with this finer dataset, the main analysis can then proceed with daily data on the full dataset.

With four-times-daily data, Figure 1 clearly shows the daily cycles in the city of Toronto before February 15th. Graphs for London and Ottawa before February 15th and the graph for Vancouver before and after February 15th are similarly clear. I find the median price change in the month before the February 15th fire is -1.05 cpl in London, -0.37 cpl in Ottawa, and -0.73 cpl in Toronto.[12]. The median price change in Vancouver in the month before the fire is -0.21 cpl.

<Figure 3 about here>

Immediately after the fire, the equilibrium in the three Ontario cities changed. Figure 3 shows the sudden cessation of cycles and reduced price volatility in London, Ottawa, and Toronto. The date of the fire, February 15th, is represented by a vertical dashed line. A similar reduction was not seen in Vancouver. The median price change after the supply shock resolved was 0.02 cpl in London, -0.06 in Ottawa and 0.001 in Toronto. In contrast, daily cycles continued in Vancouver with a median price change after March 26th of -0.45 cpl.

In terms of price volatility, the average absolute price change (every four hours throughout the day) fell from 2.91 cpl to 0.37 cpl in London, from 2.65 to 0.68 cpl in Ottawa, and from 2.42 to 0.26 cpl in Toronto from the first two weeks of February to the first two weeks of March. In Vancouver, the average absolute price change fell only modestly and insignificantly from 2.26 to 1.85 cpl over the same period.

My categorization of treatment cities agrees with Atkinson et al. (2014). It also agrees with Byrne, Leslie, and Ware (2015) which, in a study of Ontario markets, find no cycles in the treatment markets but do find weekly cycles in Windsor from August 2007 through the end of the sample.

[12] Median price changes in the daily data versus the four-times-a-day data are not directly comparable, since the former contains seven data points a week and the latter contains four data points a day.

The remaining sample markets of Calgary, Winnipeg, and Quebec City exhibit no daily or weekly price cycles over the sample period. This is evident in the daily data as prices remain sticky generally for a few days at a time. I reviewed four-times-daily subsamples for these cities to further confirm the absence of daily cycling behavior and, more importantly, the absence of any change in equilibrium immediately following the fire.

In summary, there are three treatment markets – London, Toronto, and Ottawa – all of which experienced daily cycles prior to the fire and no cycles after the fire. The daily cycles in these cities date back to 2004, and prior to that there were regular weekly to monthly cycles from at least the 1980s (Noel (2007a), Atkinson et al. (2014)).

There are six control cities – Vancouver, Calgary, Winnipeg, Windsor, Montreal, and Quebec City – whose equilibrium type did not change. The first three are western markets unconnected to the refinery and unaffected by the fire. Vancouver had daily price cycles since 2004 and continued to have them through the sample period, both before and after the fire. Calgary and Winnipeg were without cycles through the sample period and for many years prior.

The remaining three control markets are in Eastern Canada. Montreal and Quebec City are connected to the Ontario infrastructure and may have been, to a lesser extent, impacted by the shortage from the fire. However, the fire did not change the equilibrium type. Montreal experienced weekly cycles throughout the sample, both before and after the refinery fire. Quebec City did not have cycles at any time.

The last control city, Windsor, Ontario, is among the cities located near Nanticoke. Windsor experienced weekly to monthly cycles since the 1980s, and weekly cycles leading up to the fire. The weekly cycles were disrupted by the fire, but unlike other cities, they restarted on a weekly to biweekly basis once the refinery came back fully online.

Together, the sample cities represent a good mix of equilibrium types and a geographically well balanced set of treatment and control cities – three treatment cities, three control cities in the West, three control cities in the East. Treatment and control cities are also evenly distributed in terms of population ranking. Treatment cities rank 1^{st}, 4^{th}, and 8^{th}; control cities rank 2^{nd}, 3^{rd}, 5^{th} through 7^{th}, and 9^{th}.

The technique is a difference-in-differences estimation on a panel of nine cities surrounding the refinery fire shock. The baseline estimating equation is given by:

$$\begin{aligned}
RETAIL_{mt} \quad = \quad & \beta_0 + \beta_1 ADJUSTMENT_PERIOD_t + \beta_2 POST_SHOCK_t \\
& + \beta_3 TREATMENT_m * ADJUSTMENT_PERIOD_t \\
& + \beta_4 TREATMENT_m * POST_SHOCK_t \\
& + \beta_5 RACK_{mt} + \sum_{m=2}^{9} \phi_m + \sum_{j=2}^{12} \xi_j + \varepsilon_{mt}
\end{aligned} \tag{4}$$

where $RETAIL_{mt}$ is the price of regular grade gasoline for market m in period t. $RACK_{mt}$ is the tax-adjusted wholesale price and the ϕ_m and ξ_j are market and monthly fixed effects, respectively. The ε_{mt} are normally distributed error terms, correlated within-markets. I cluster standard errors by city in all specifications to account for this dependence.

In addition to the baseline model, I consider a wide range of alternate specifications to check the robustness of the results. These include models with different subsets of control cities and different subsets of treatment cities, models using different event windows, models that account for short run dynamics such as dynamic least squares (DOLS) and vector autoregressive error correction models (ECM), models that explicitly control for cycle position in cycling cities, and models using different levels of data aggregation guarding against serial correlation concerns. These models are discussed more fully as they are presented below. I also conduct a series of falsification exercises that support the main conclusion.

The standard difference-in-differences approach is to include a treatment market variable, a post-shock period variable, and the interaction of the two is the variable of interest. However, it is imprudent in this setting to include the weeks immediately following the fire in the post-shock period. Even though cycles ceased within a week of the fire, this period involved a supply shortage that temporarily drove up rack prices. While I am not interested in rack price changes per se, if the increase in rack prices did not fully reflect the wholesale shortage, it may also affect retail margins and retail prices after controlling for rack prices.

As a result, I consider three different dates for the start of the post-shock period. The first date is based on the date the Nanticoke and other southern Ontario rack prices returned to pre-fire levels following the fire. Prior to February 15, 2007, the Nanticoke rack was priced at a 0.54 cpl discount relative to non-Ontario cities. It rose and peaked at a 4.57 cpl premium on March 8th. On March 15, 2007, Imperial Oil, the operator of the refinery, announced that refinery repairs were complete and the refinery was back in full operation. Rack prices had already been falling and on March 26th, Nanticoke and other southern Ontario racks returned to the discount they had on February 14th. Therefore, I use March 26th as the primary date for the start of the post-shock period, and $POST\_SHOCK_t$ takes on the value of one after March 26th.[13]

I consider two alternate dates as well. Although affected racks returned to pre-fire levels by March 26th, it may be that retail margins took longer to recover, assuming retail margins were affected to begin with. The lag between wholesale price movements and retail price movements has been well studied and can potentially last a few weeks (Borenstein, Cameron, and Gilbert (1997), Lewis and Noel (2011)). To be conservative, I use April 30, 2007, two and a half months after the fire and five weeks after rack prices returned to normal, as an alternate date for the start of the post-shock period. To be even more conservative, I also use a second alternate date of June 15, 2007, four months after the fire and eleven and a half weeks after wholesale prices renormalized. The main purpose of the later date is to show the margin changes following the fire remained in place in the long run, long after any short run dynamics have played out. I find very similar results for all three dates and all intermediate dates.

In terms of the intermediate adjustment period, one option is to drop it from the analysis altogether. However, since this data can still identify other coefficients (e.g. rack), I maintain it in the analysis. I create an indicator variable $ADJUSTMENT\_PERIOD_t$ equal to one during the adjustment period and allow it to vary across treatment and control markets (using an $ADJUSTMENT\_PERIOD_t * POST\_SHOCK_t$ interaction). Results are similar including this period and these controls, or excluding the period entirely.

I set the indicator variable $TREATMENT_m$ equal to one for the cities of London, Ottawa, and Toronto. I include market-specific fixed effects in all specifications, obviating the need for a treatment main effect. The coefficient of interest is β_4, the coefficient on the $TREATMENT_m * POST\_SHOCK_t$ interaction term, and represents the difference in the change in price in treatment

[13] Nanticoke and other southern Ontario racks were highly correlated, with an average correlation coefficient of 0.98.

markets relative to control markets pre- and post-shock.

I also estimate an isomorphic set of margin regressions, with $MARGIN_{mt}$ as the dependent variable, where the retail-rack spread is used as a proxy for margins.[14] The margin model is similar to Equation (4) but with the restriction $\beta_5 = 1$.

Before turning to the main results, I perform a series of diagnostic checks. First, I test for unit roots in the retail price, rack price, and margin time series in each city both before and after the known February 15th break. Phillips-Perron unit root tests cannot reject the null hypothesis of a unit root in rack and retail prices at the 5% level in any of the markets (and in only five of nine at the 10% level). However, they do reject the null hypothesis of a unit root in retail-rack margins in every market at better than the 1% level. Augmented Dickey-Fuller unit root tests agree. Engle Granger test and the Johansen test of cointegration also agree that prices and racks are cointegrated in each market, with a rack price coefficient of 0.99, very close to one.[15] Finally, I estimate treatment-group-specific time trend regressions for the pre-shock period and confirm the trends are insignificantly different across treatment and control groups (with a t-statistic of 0.19), as needed in difference-in-differences models.

4 Results

In short, I find that the elimination of Edgeworth price cycles led to higher retail prices and margins in the affected markets. In other words, the presence of Edgeworth price cycles in these markets caused a decrease in retail prices and retail margins. The results are robust to a wide variety of alternate specifications.

In the top panel of Table 2, I report baseline results on the effects of the cessation of cycles on prices. In the bottom panel, I report its effects on margins. I report three specifications for each panel, corresponding to the alternate dates for the start of the post-shock period – March 26th, April 30th, and June 15th. The coefficient of interest is the $TREATMENT_m * POST\_SHOCK_t$ interaction term. All regressions include market and monthly dummies, and standard errors are

[14] Since I estimate differences in the changes in margins across treatment and control cities, it is not important that the retail-rack spread is only a proxy for margins.

[15] The Engle Granger test is based on an Augmented Dickey Fuller test on the residuals of the known cointegrating regression, but using further adjusted critical values from Phillips and Ouliaris (1990).

clustered by city to account for dependence within panels.

The main result is that the cessation of cycles led to higher prices and margins, with estimates ranging in a small band between one and one-and-a-half cpl, and all highly statistically significant. For the March 26th post-shock start date, I find retail prices in the treatment markets rose by 1.31 cpl relative to control markets, significant at the 1% level. The estimates for the April 30th and June 15th post-shock start dates are 1.09 cpl and 1.06 cpl respectively, also both significant at the 1% level. The latter result in particular shows the price increase did not go away even beyond any realistic time frame for short run dynamics to play out following the refinery's recovery. It is also consistent with the result that margins were stationary in each city before and then after the fire. The estimates correspond to economically significant price increases of 4.5, 3.8, and 3.7 U.S. cents per gallon, and support the conclusion that cycles had led to lower prices.

In terms of the other coefficients, the coefficient on $RACK_{mt}$ is 0.99 cpl and insignificantly different than one in all specifications, consistent with near complete long run pass-through of rack prices into retail prices. The coefficient on the $TREATMENT_m * ADJUSTMENT\_PERIOD_t$ interaction confirms the need to separately control for the adjustment period. The coefficient shows that prices in the treatment cities rose 1.9 to 2.5 cpl more than control cities right after the fire before rebounding back.

The bottom panel reports difference-in-differences estimates on margins. Margins in the treatment cities increased between 1.06 to 1.32 cpl above those in control cities, depending on the choice of the post-shock start date. The effects are very similar to those for prices, not surprisingly given the rack price coefficient is so close to unity. In percentage terms, the change is substantial – ranging from 19% to 23% of the mean retail-rack margin as measured in the sample.

I now turn to alternate specifications. They are reported in Tables 3 and 4, with price effects in the top panel and margin effects in the bottom. To conserve space, I report only the coefficient of interest – the $TREATMENT_m * POST\_SHOCK_t$ interaction – along with its standard error, for each of the alternate specifications. The specification relevant for each cell is described in the corresponding row using the post-shock start date from the corresponding column.[16] In all, Tables 3 and 4 present results from seventy-two different specifications in addition to the baseline six.

[16] The full set of coefficients from each specification is available upon request.

First, I test whether the results are sensitive to the set of control markets chosen. In the first row of each panel, I restrict the control markets to be only those in the sample from Western Canada – Vancouver, Calgary, and Winnipeg. I find the results are very close to the baseline models. Prices and margins in the treatment markets increased by between 1.10 and 1.36 cpl over control cities after the fire, statistically significant at better than the 1.2% level.

In the second row, I restrict the control markets to be only those in the sample in Eastern Canada – Windsor, Montreal, and Quebec City. The point estimates are almost identical to the baseline model, and to the Western Cities as Controls model. Prices and margins increased in treatment cities by between 1.12 to 1.39 cpl across post-shock start dates. Because the data for the latter two cities begin only in June 2006, the clustered standard errors are about double those in the first row. Nonetheless, the result is still significant (at the 1% and 10% levels) for two of the three post-shock start dates, the first date (March 26th) and the last date (June 15th), while very close on the other.

Another robustness check treats the control city of Windsor differently. Windsor is located in southern Ontario and is nearby the Nanticoke refinery. Its cycles stopped along with those in the treatment cities but, unlike the treatment cities, its cycles restarted after the fire. However, upon closer inspection, it is revealed that the nature of those cycles did in fact change with the restart. Prior to the fire, cycles were roughly weekly with an average duration of 8.00 days and an average amplitude of 7.34 cents per liter. When the new cycles restarted after the fire, they were more often biweekly instead of weekly and had a much reduced amplitude. The average duration rose to 10.16 days and the amplitude fell to 4.89 cpl, both statistically significant changes.

Edgeworth price cycles can take on many shapes, and shorter amplitude and longer period cycles are often associated with weaker competition (Noel (2007a, 2008)). If the fire shocked the intensity of cycles in Windsor, even if it did not shock the equilibrium type, then Windsor may have experienced some degree of treatment, understating effects.

To address this, in the third row specification "Using Only Quebec and Western Controls", I exclude Windsor as a control market. I find the coefficient on $TREATMENT_m * POST\_SHOCK_t$ in the price regressions now ranges from 1.11 to 1.39, very similar to previous results, and statistically significant at the 5% level or better. Margin results are also similar.

For completeness, I also compare the treatment cities to Windsor alone (not shown). I still find significant treatment effects, from 0.74 to 0.87 cpl, and significant at the 2% level. The slightly lower estimates suggest that Windsor did experience a small degree of treatment from the fire by slowing down the cycle intensity. However, the price increase due to the cessation of the cycles still significantly exceeded any price increase due to the slowing of the cycles. I conclude the treatment effect is robust to the inclusion or exclusion of Windsor.

Another robustness check takes a closer look at pricing and margin patterns in the treatment cities after the fire. In their examination of Toronto, Atkinson et al. (2014) find that when cycles stopped in Toronto, margins became almost constant across stations and over time at 5 cpl. This is an interesting result, as one may have expected margins in non-cycling markets to show more variation, albeit only a little since gasoline is relatively homogenous.

The potential concern is as follows. If the almost constant margins in Toronto after the fire were a sign of especially or unusually weak competition, and if the other treatment cities similarly switched to constant margins like Toronto, and finally if constant margins are not representative of what would normally be expected when cycles stop elsewhere, the treatment effects estimated here may not generalize well. The same concern holds if the constant margins in Toronto – which is equivalent to immediate and symmetric pass-through of rack price changes into retail prices – were a sign of unusually strong competition and not representative of what might occur elsewhere if cycles stopped.[17]

So the first step is to check which treatment cities switched to a constant margin equilibrium and which, if any, switched to an equilibrium that maintained some degree of margin variance. If all treatment cities switched to constant margins after the fire, the results would be applicable to such cities but may or may not be applicable to other situations. If instead, both constant and variable margin patterns are found in the treatment group, then one can estimate price and margin effects of the fire for each group separately and see if the effects are similar or different.

I plot retail-rack margins for the three treatment cities – Toronto, Ottawa, and London – over

[17] Immediate and complete passthrough is the extreme case of cost-based pricing. Noel (2007a) and Byrne, Leslie and Ware (2015) associate more (typical) cost-based pricing with greater competition within the universe of non-cycling markets. Interestingly, the opposite of immediate and complete passthrough – slow and asymmetric passthrough – has been itself criticized as reflective of less competition (Borenstein, Cameron, and Gilbert (1997), Deltas (2008)).

time in Figure 4. The date of the fire, February 15th, is again represented by a vertical dashed line. As an example of a control city, I show margins for Vancouver, which experienced daily cycles both before and after.

The constant margins of 5 cpl in Toronto are clearly evident in the top left panel. However, and interestingly, this pattern is not universal across treatment cities. In Ottawa, volatility decreased (which can be expected when volatile prices are replaced with more stable prices) and while margins could arguably resemble something close to a constant pattern up until the June post-shock date, margins after the June post-shock date were clearly not constant. The case is even clearer in London. In London, volatility decreased relatively little after the fire and there is no sense of a constant margin at any time.

So both types of equilibria – constant and non-constant margins – are present in the post-shock world. Accordingly, I estimate two more separate treatment effects. In the fourth row of each panel of Table 3, I restrict the treatment group to Toronto only, with its clear constant margins, and, in the fifth row, I restrict the treatment group to London and Ottawa only. As the table shows, I find significant treatment effects in both cases. The coefficient on $TREATMENT_m * POST\_SHOCK_t$ in the "Toronto Only vs. Controls" price specifications ranges from 1.39 to 1.60, statistically significant in each case at the 1% level or better. The coefficient in the "Ottawa & London vs. Controls" price specification ranges from 0.89 to 1.15, statistically significant in each case at the 1.5% level or better. Notably, this latter result holds even for the post-June start date where margins are clearly non-constant in both cities. Coefficients in the margin regressions are very similar. Taking this further, specifications isolating Ottawa as the sole treatment city and then isolating London as the sole treatment city yield very similar results to each other and to the combination. In other words, the main result holds regardless of the observed post-equilibria pattern. When cycles stop and are replaced with constant margins, prices and margins rise. When cycles stop and are replaced with more variable margins, prices and margins still rise.[18].

The slightly higher effects in Toronto may appear to suggest that the constant margin pattern

[18] A third type of pattern, "sticky pricing" in the nomenclature of Noel (2007a), in which prices routinely remain fixed for weeks or months at a time, is not observed in the sample. Sticky pricing has typically been observed in smaller towns with fewer stations and in Atlantic Canada (Noel (2007a), Byrne and Leslie (2013)). They have not been seen in large Ontario markets since at least the 1980s. Prices exhibit relatively little connection to costs in these markets and prices tend to be among the highest.

was less competitive than the variable margin pattern. However, caution is warranted. It could instead be that the cycles in Toronto were especially competitive before the fire. Noel (2008) finds that cycles of varying intensities lead to different price and margin levels and it is known Toronto had especially intense cycles relative to others (Noel (2007a)). Margins in Toronto were also the lowest in the sample before the fire.

The robustness of results to various subsets of control and treatment cities can also be seen informally with a simple comparison of margin changes. The average increase in margins in the treatment cities is 2.6 cpl, ranging from 2.3 cpl in London to 2.9 in Toronto. The average increase in margins in control cities, in contrast, is just 1.2 cpl and none exceed 2 cpl. In other words, the margin increase in every treatment city is greater than the margin increase in any control city and therefore no combination of treatments and controls can reverse the result.

Another check narrows the length of the estimation window. In the sixth row, the "Reduced One-Year Event Window" specification, I use data only from six months prior to the fire and six months after the start of the post-shock period, thus disregarding about half the data. Doing so more than doubles the standard errors, but the treatment effects remain positive and significant at better than the 10% level. Price and margin increases were 1.41 to 1.83 higher in treatment cities than in control cities. Other estimation windows produce expected results and agree with the main conclusion.

Next, I turn to short run dynamics. Although the focus is on estimating long run changes in prices and margins due to the cessation of cycles, there are two motivations for adding short run dynamics.

First, prices and racks are cointegrated, so short run dynamics contain additional information that can be used to estimate standard errors corrected for superconsistency in the cointegrating regression (the long run retail-rack relationship). While I wish to err on the side of being conservative, it is instructive to see how the estimated standard errors and potentially the estimated treatment effects change. Superconsistency-corrected standard errors are estimated using the dynamic ordinary least squares (DOLS) specification given by Stock and Watson (1993),which simply adds lagged and leading rack price changes to Equation 4:

$$
\begin{aligned}
RETAIL_{mt} &= \beta_0 + \beta_1 ADJUSTMENT\_PERIOD_t + \beta_2 POST\_SHOCK_t \\
&+ \beta_3 TREATMENT_m * ADJUSTMENT\_PERIOD_t \\
&+ \beta_4 TREATMENT_m * POST\_SHOCK_t \\
&+ \beta_5 RACK_{mt} + \sum_{m=2}^{9} \phi_m + \sum_{j=2}^{12} \xi_j \\
&+ \sum_{s=-S}^{S} \gamma_s \Delta RACK_{m,t+s} + \varepsilon_{mt}
\end{aligned} \tag{5}
$$

and calculates standard errors using a heteroskedasticity and autocorrelation consistent variance covariance matrix (e.g. Newey West). The margin equation is the same except with $\beta_5 = 1$. The common lag and lead length S was chosen by the testing down method.[19].

The results are reported in the first row of each panel of Table 4. The coefficient on $TREATMENT_m * POST\_SHOCK_t$ in the price regressions ranges from 1.26 to 1.49 cpl, similar to previous results. Margins regressions yield very similar results as well. Standard errors estimates are indeed modestly lower than in the baseline model, but as coefficients were already statistically significantly at better than the 1% level, the main conclusions do not change.

The second motivation for including short run dynamics is that the fire temporarily shocked rack prices in the treatment cities, and rack prices tend to pass into retail prices with a lag. If the short run dynamics took long enough to play out, the effects of the shock could in principle infect the post-shock period. Consequently I estimate a vector autoregressive error correction model in the spirit of Engle and Granger (1987) of the general form:

$$
\begin{aligned}
\Delta RETAIL_{mt} &= \sum_{s=0}^{S} \gamma_s^+ \Delta RACK_{m,t-s}^+ + \sum_{s=0}^{S} \gamma_s^- \Delta RACK_{m,t-s}^- \\
&+ \sum_{r=1}^{R} \delta_r^+ \Delta RETAIL_{m,t-r}^+ + \sum_{r=1}^{R} \delta_r^- \Delta RETAIL_{m,t-r}^- \\
&+ \theta^+ z_{m,t-1}^+ + \theta^- z_{m,t-1}^- + \varepsilon_{mt}
\end{aligned} \tag{6}
$$

[19] In the reported specifications, I use S=15. Results are similar for other values. Leads are almost never significant at any lead length.

where

$$
\begin{aligned}
z_{mt} \;\; = \;\; & RETAIL_{mt} \; - \beta_1 ADJUSTMENT\_PERIOD_t - \beta_2 POST\_SHOCK_t \\
& -\beta_3 TREATMENT_m * ADJUSTMENT\_PERIOD_t \\
& -\beta_4 TREATMENT_m * POST\_SHOCK_t - \beta_5 RACK_{mt} - \sum_{m=2}^{9} \phi_m - \sum_{j=2}^{12} \xi_j \qquad (7)
\end{aligned}
$$

Let $A^+ = max(0, A)$ and $A^- = min(0, -A)$ for $A = \{RACK_{mt},\ RETAIL_{mt},\ z_{mt}\}$. I perform both one-step (Banerjee et al. (1993), Borenstein, Cameron, and Gilbert (1997)) and two-step (Engle and Granger (1987), Bachmeier and Griffin (2003)) versions of the model. In the one step model, Equation 7 is substituted directly into Equation 6 and then estimated.[20] In the two-step version, Equation 4 is estimated in the first stage, and the first stage residual z_{mt} is taken as known (instead of estimated) due to superconsistency and substituted into Equation 6 for the second stage estimation of the short run dynamics.

I begin with the two-step asymmetric model. By its construction, the estimated coefficient on the $TREATMENT_m * POST\_SHOCK_t$ variable is identical to that reported in Table 2, since the first step is nothing more than estimating Equation 4 again. Thus the two-step ECM replicates the main conclusion and I do not re-report the treatment effects here. It is worth noting that the short run dynamics coefficients from this model show that asymmetric pass-through is very fast and not very asymmetric. The coefficient $\gamma_1^+ = 0.62$ (first day percentage response after an increase) and coefficient $\gamma_1^- = 0.49$ (first day percentage response after a decrease) show that half or more of pass-through is complete within a day. I find 75% is complete in four days in either direction, 90% in seven days in either direction, and virtually all in nine days. This is broadly consistent with other studies including Eckert (2003), Noel (2009), and Byrne, Leslie, and Ware (2015).

It suggests lingering effects from short run dynamics in the retail-rack relationship could not meaningfully spill into post-shock period and drive results. The first post-shock start date is March 26th, the date rack prices had returned to pre-shock levels, relative to control racks. The second treatment start date of April 30th is almost a month after pass-through would have long been

[20] The one step version requires $z^+_{mt} = z^-_{mt}$ but maintains direction specific γ^+, γ^-, δ^+, and δ^- in the estimation. In both, I use $S = 40$ and $R = 15$, though results are extremely similar with smaller lag lengths as well.

complete. The third treatment start date of June 15th is four months beyond the fire and so far out that pass-through lags cannot reasonably play a role.

This can also be tested directly. Since the two-step method only replicates Table 2, I perform a one-step version which estimates Equations 6 and 7 together – the "Error Correction Model" in the second row of each panel of Table 4. I find the coefficient on the $TREATMENT_m * POST\_SHOCK_t$ variable ranges from 1.08 to 1.29 cpl, statistically significant at the 1% level. These are very similar to the results previously reported in Table 2 and the alternate specifications of Table 3. I also perform a margins version of the error correction model constraining $\beta_5 = 1$ and again find similar estimates to before, ranging from 1.09 to 1.33 cpl. I conclude short run dynamics do not have a meaningful impact on estimates of the long run impact of cycles on margins and prices.

A different robustness check takes advantage of the deterministic nature of prices along the cycle and adds cycle position as a control variable. While the longest cycles in the data (in Windsor) are only a few weeks long and unlikely to impact the long run difference-in-differences estimates, controlling for cycle position can potentially improve accuracy. For daily cycles, the issue is moot because data is aggregated to the daily level for use in the regressions already. For weekly cycles (in Montreal and Windsor), one can control for the position in the cycle by counting the number of days since the last peak (Noel (2012), Noel and Chu (2015)). I do this in the third row specification in each panel of Table 4, "Include Cycle Position Measure", and find the coefficient on $TREATMENT_m * POST\_SHOCK_t$ ranges from 1.02 to 1.25 cpl, statistically significant at the 1% level, and again very similar to the baseline model.

Another concern about difference-in-differences models generally is that serially correlated data in a long panel can potentially lead to downward biased standard error estimates and incorrect rejection of the null. Bertrand, Duflo, and Mullainathan (2004) show in Monte Carlo simulations that standard OLS implementation with no standard error corrections can reject a null hypothesis of no effect (at the 5% level) almost 50% time when there is really is no effect.

I address this in a number of ways. Most importantly, in all specifications, I use an arbitrary variance-covariance matrix to estimate standard errors, i.e. clustering, to account for correlation in the error term. Bertrand et al. show this adjustment largely cures the issue when there are many clusters and does well, though not perfectly, for smaller numbers of clusters. With ten clusters (close to the nine in this study), the rejection rate falls to 8% (instead of 50%) when the null is true and the 5% significance level is used.

Since the potential for over-rejection remains a concern, I implement a technique that Bertrand et al. find works best with small numbers of clusters in reducing Type I error. They suggest collapsing or aggregating up the data to a coarser unit of time, essentially averaging serially correlated observations into a single observation, and removing much of the time dimension from the analysis. The approach reduces Type I error to the correct level, even with few clusters, but comes at the cost of substantially increasing Type II error (not rejecting a null hypothesis that is actually false). In other words, it can fail to detect effects that are really there, and is thus a trade-off.

To implement this method, I perform three more sets of specifications that collapse the data to different levels. Since serial correlation is likely to be most strong within a cycle, I first collapse the data to the weekly level, so both daily and weekly cycles are reduced to a single weekly average. This reduces the total number of observations by 86% of the original, removing potentially problematic serial correlation but at the cost of a higher Type II error. I then collapse the data to the monthly level, eliminating any within-month serial correlation but losing 96.8% of the original number of observations. Finally, I take this to the extreme and collapse the data from over 6000 data points to just 27 data points – one for each city for the pre-shock, adjustment period, and post-adjustment periods respectively. This is exceptionally demanding of the data as it reduces the number of observations by 99.5% and potentially significantly inflates Type II error since both good and potentially bad variation is being washed away into single averages.[21]

Impressively, the results of all three specifications, even that with just 27 observations, are very similar to the previous results and all maintain statistical significance. I report results from the "Collapse Data to Monthly Level" and "Collapse Data to Pre/Post Level" specifications in the fourth and fifth rows of each panel of Table 4. The results for the weekly level specification are similar to that of the monthly level specification and are not reported.

[21] In their simulations, Bertrand et al. reject the null hypothesis of no effect only 6% of the time when the alternative hypothesis (of a 2% effect on the left hand side variable) is actually true. That is, statistical power is just 6% and Type II error is high at 94%. Fortunately, power increases and Type II error falls the greater is the true effect (under the alternative hypothesis) above 2%.

In the monthly level regressions, the coefficient on $TREATMENT_m * POST\_SHOCK_t$ in the price regressions ranges from 1.11 to 1.32 cpl, similar to previous estimates, and significant at the 1% level or better. The corresponding coefficients in the margin regressions range from 1.12 to 1.31 cpl, also significant at the 1% level. In the pre/post level regressions that use just three data points per city, 27 in all, the coefficient on the $TREATMENT_m * POST\_SHOCK_t$ variable ranges from 1.15 to 1.59 cpl, and is still significant at the 5% level or better. This is an exceptional result. The corresponding coefficients in the margin regressions range from 1.08 to 1.39 cpl. I conclude that serial correlation bias and the potential for excessive Type I error is not affecting results in any meaningful way.

There is a second way to gauge the potential for an incorrect rejection from serial correlation bias. Bertrand et al. demonstrate the potential for over-rejection by imagining a series of "placebo laws" at random dates and performing difference-in-differences regressions around those dates in search of an effect that is not really there. I can do the same. I conduct a series of falsification exercises in which I imagine fake events taking place on random dates and search for a significant effect where I would not expect to find one. If this happens more than 5% of the time (at the 5% significance level), then unaccounted-for serial correlation in the data may be an issue.[22]

I divide the data into two intervals, before the fire and after the adjustment period is over, so as not to pick up any effect of the fire, and randomly select twenty event dates in each interval. I rerun the full analysis, with clustered standard errors, around these dates in the relevant interval. In every falsification exercise, I do not reject the null hypothesis of no effect when I do not expect to see an effect. In other words, my rejection rate is 0% (instead of the allowable 5%), and I conclude serial correlation is not leading to an over-rejection problem in this context.[23]

The main contribution of this article has been to estimate the causal effect of Edgeworth price cycles on prices and margins. In spite of the large literature on cycles and varying hypotheses about their competitiveness, the question of whether they lead to higher or lower prices has been

[22] The discussion assumes there were not any "real" events that would cause meaningful and discrete changes in prices and margins over time. Given this, false rejections are merely picking up the effect of unaccounted-for serial correlation of small, random shocks (see Bertrand et al.). To the extent there actually could have been real, confounding nuisance events (I find there are not), these would be a concern in and of themselves, as discussed below.

[23] This also shows there is no evidence of confounding nuisance events that could infect results. I also performed other adjustments considered by Bertrand et al., including parametric AR(1) adjustments and Newey West standard error adjusted regressions. Results are similar in magnitude and significance to the baseline results.

elusive. The identification of a sudden and exogenous shock to equilibrium types, and a difference-in-differences framework wrapped around that event to control for trends and common shocks, is new to the literature and important is that regard.

To contrast the difference-in-differences model to a straight before-and-after comparison, I consider a regression discontinuity type regression in the last row of Table 4. According to this specification, which does not use control markets, post-shock prices rose between 2.30 and 2.58 cpl in the treatment cities after the fire and margin increases were similar. (The coefficient is even higher if one ignores the need to exclude an adjustment period.) This result, as it turns out, is biased upward.

The full difference-in-differences model cuts the estimate in half. Ex ante, it was just as reasonable that the difference-in-differences estimation may have reduced, eliminated, or even reversed the regression discontinuity estimate and yielded a negative price change when cycles stopped. That is why both the use of an exogenous shock to the equilibrium (the fire) and a panel approach to control for common unobserved shocks across markets is important. As it turns out, margins were not only increasing in the treatment cities right around the time of the fire but also in the other cities as well, confounding a causal interpretation from the regression discontinuity estimates alone.

Similarly, a simple cross-sectional comparison before the fire would show prices in treatment cities with cycles were 3.4 cpl lower than in cities without. Again, such estimates are prone to bias and are easily dismissed as unreliable absent a difference-in-differences framework.

It is worth noting that the results of this article, even with a difference-in-differences framework, may still modestly understate the benefits of Edgeworth price cycles. Like virtually all studies of retail gasoline markets, this study relies on comparisons of average prices as recorded in publicly available data.[24] Noel (2012) and Noel and Chu (2015) show that, where price cycles are present and predictable, consumers may be able to time purchases to periods near the troughs.[25] In markets

[24] Matched price-quantity data at the necessary frequently are almost never available. Two rare exceptions using matched data are Barron et al. (2008) and Wang (2009a), who obtained proprietary data (and in one case, were able to control a firm's prices).

[25] This suggests a hidden advantage of Edgeworth price cycles, even when average (unweighted) prices with or without cycles are the same. The Australian Competition and Consumer Commission (ACCC, 2007) reports that 75% of Australian consumers in the four major eastern cities were aware of the weekly cycles and 74% of those knew which day of the week was the low price day. Foros & Steen (2008) report that about one third of Norwegian consumers were aware of the weekly cycle in Norway.

without cycles, this option does not exist. To the extent consumers exploit this ability and buy more on days with lower prices than higher prices, average prices under cycles are even lower.

In summary, the results of this section show that Edgeworth price cycles prior to the February 15th fire resulted in lower prices and margins for markets with cycles. When the cycles suddenly stopped and more stable pricing returned, retail prices and margins both went higher. The result is robust to a wide variety of static and dynamic models. It is also robust to whether constant margins or variable margins emerged in the new non-cycling equilibrium. The change in margins was especially substantial in percentage terms.

5 Discussion

This is the first study to exploit a sharp natural experiment and a panel setting to isolate the effects of Edgeworth price cycles on retail prices and margins. The results qualitatively agree with results from previous studies that make rough comparisons of prices across cities with and without cycles (Noel (2002), Doyle et al. (2010)), or compared prices with and without cycles over a long time horizon (Zimmerman et al. (2013)). They are also consistent with indirect evidence under the Edgeworth price cycle theory that cycles are more likely with more price aggressive independent firms (Eckert (2003), Noel (2007b), Doyle et al. (2010), and Lewis (2012)). They are not as consistent with a collusion-based or anti-competitive view of cycles which would presumably raise prices (Foros & Steen (2013), de Roos and Katayama (2013)).

The sudden discontinuity in the equilibrium pattern is the key to unlock the causal effect of cycles. Prices and margins changed almost instantaneously and apples to apples comparisons can be made within months of each other, controlling for shocks or trends common to both affected and unaffected cities. One remaining concern might be that the fire, in addition to shocking the equilibrium types, also suddenly and discretely shocked market structure in the treatment cities at the same time. However, the evidence shows this is not the case. While in the early weeks following the fire, stations were put on allocation and some ran out of fuel (enjoying higher margins on the fuel they did sell), the use of an adjustment period of one to four months long removes the allocation effect. Once the refinery was back to full operation and rack prices fully recovered by late

March, there is no evidence of unusual exit. Atkinson et al. (2014) contains an excellent discussion. Five-firm concentration ratios in the London, Toronto, and Ottawa areas were virtually unchanged at 0.01, -0.01, and -0.03 between 2006 and 2009.[26] There was also no merger or other significant activity at that time. Further, a search of Toronto newspaper archives reveals many stories about the fire but none about widespread, or in fact any, unusual station closures in the direct aftermath. Finally, given the vast majority of (non-supermarket) stations are branded stations with lessee dealers that have supply contracts of six or ten years (with financial penalties for early termination), a disruption lasting a month that raised margins was unlikely to trigger sudden mass exit.

The results of the study have implications for competition policy. Edgeworth price cycles have been controversial. They have inspired investigations by federal competition authorities in a number of countries (e.g., the U.S., Australia, Germany, Norway) and have led, in one case, to a form of price regulation. Investigations tend to focus on the rapid sequence of price increases and, with no obvious contemporaneous changes in costs to justify them, the concern is that price increases are collusive in nature (Foros and Steen (2013), De Roos and Katayama (2013)). However, less attention is given to the undercutting phase or to the question of whether or not the cycle, taken as a whole, ever led to higher prices in the first place. Lewis and Noel (2011) show that cycle peaks are higher and cycle troughs lower than constant prices in observationally similar cities without cycles, so the narrow focus on just the "bad" half of the cycle process seems incomplete.

In a recent German investigation, the competition authority investigated cycles and, citing strong parallel price movements along them, declared the top five oil companies collectively dominant (Bundeskartellamt (2011)). The declaration affects the handling of future antitrust and merger matters in the industry. Recent Norwegian investigations looked at the cycles but did not uncover evidence of collusion (Konkurransetilsynet (2011)). At the time of writing, an investigation of price cycles in Australia is underway. In the United States, in 2011, the Federal Trade Commission released a report addressing the cycles and advocating no action pending further learning on the topic (Federal Trade Commission (2011)). This article contributes to that learning.

Other jurisdictions have considered price regulation to offset the perceived negative effects of

[26] These data, collected by Kent marketing, are based on volume shares of participating stations, and do not include large supermarkets.

the cycle. The FuelWatch program implemented in Western Australia in 2001 requires firms to provide 24 hour pre-notification of price changes and limits them to, at most, one change per day. In 2007, the Australian government considered extending the program nationwide, leading Harding (2008) to coin the phrase "FoolWatch". A similar bill requiring 72 hour notice of price increases failed to pass in Canada in 2007 (Noel (2012)). As long as Edgeworth price cycles continue to spark consumer concern, legislative and regulatory activity is likely to continue.

It is important to note that while I find Edgeworth price cycles in the sample markets lowered prices and margins, even if cycles had no significant effects of cycles on prices or margins at all, the policy implications would still be largely the same. The evidence suggests that Edgeworth price cycles are not harmful to consumers and are not a public policy problem in need of fixing.

6 Conclusion

In this article, I exploit a sharp natural experiment to identify the causal effects of price cycles on retail price levels and margins. I focus on a refinery fire which took place at the Nanticoke refinery in southern Ontario on February 15, 2007. The fire caused the retail price cycles that had persisted in several nearby cities for decades to suddenly stop. In other cities further afield, the fire did not change the equilibrium type – cities with cycles continued to have cycles, and cities without continued without. I find the cessation of Edgeworth price cycles, all else equal, led to a price increase of 1.06 to 1.31 Canadian cents per liter, or 3.7 to 4.5 U.S. cents per gallon. In other words, the presence of Edgeworth price cycles lowered prices.

A key advantage of this study over past efforts is its sharp natural experiment design and short panel setting to causally identify effects. The challenge in the literature has been that cycles tend to persist for decades, often predating the start of sufficiently fine data able to observe them. As a result, most previous studies relied on cross-sectional comparisons across cities and all were subject to concerns of potential omitted variables bias, especially with respect to market structure. In this article, identification is enhanced because I observe a clear discontinuity in the permanent equilibrium type in the treatment cities. I compare prices in the same cities under the same market conditions, once with regular cycles and once permanently without, and all within months of each other. I use a panel that also controls for changes in prices and margins in unaffected cities to remove potential confounding factors and zero in on the causal effect of interest.

The results have important policy implications. Competition authorities are generally wary of the cycles. But absent Edgeworth price cycles as a whole actually being harmful to consumers, there is no consumer gain to be had by attempting to eliminate them or regulate them away. In fact, such efforts run the risk of doing more harm than good.

7 References

Australian Competition and Consumer Commission (ACCC) (2007), 'Petrol Prices and Australian Consumers – Report of the ACCC Inquiry into the Price of Unleaded Petrol', ACCC Report, 1-280.

Allvine, F. and J. Patterson (1974), Highway Robbery: An Analysis of the Gasoline Crisis, Bloomington, IN: Indiana University Press.

Atkinson, B. (2008), 'On Retail Gasoline Pricing Websites: Potential Sample Selection Biases and their Implications for Empirical Research', *Review of Industrial Organization*, 33 (2), 161-175.

Atkinson, B. (2009), 'Retail Gasoline Price Cycles: Evidence from Guelph, Ontario Using Bi-Hourly, Station-Specific Retail Price Data', *Energy Journal*, 30 (1), 85-109.

Atkinson, B., A. Eckert and D.S. West (2014), 'Daily Price Cycles and Constant Margins: Recent Events in Canadian Gasoline Retailing', *Energy Journal*, 35 (3), 47-69.

Bachmeier and Griffin (2003), 'New Evidence on Asymmetric Gasoline Price Responses', *Review of Economics and Statistics*, 85 (3), 772-776.

Banerjee, A, J. Dolado, J. Galbraith, and D.F. Hendry (1993), Cointegration, Error Correction, and the Econometric Analysis of Non-Stationary Data, Oxford: Oxford University Press.

Barron, J.M., J.R. Umbeck and G. R. Waddell (2008), 'Consumer and Competitor Reactions: Evidence from a Field Experiment', *International Journal of Industrial Organization*, 26 (2), 517-531.

Bertrand, M., E. Duflo and S. Mullainathan (2004), 'How Much Should We Trust Differences-in-Differences Estimates?', *Quarterly Journal of Economics*, 119 (1), 249-275

Bloch, H. and N. Wills-Johnson (2010), 'A Simple Spatial Model for Edgeworth Cycles', *Eco-*

nomics Letters, 108 (3), 334-336.

Borenstein, S., A.C. Cameron, and R. Gilbert (1997), 'Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?', *Quarterly Journal of Economics*, 112 (1), 305-339.

Bundeskartellamt (German Competition Authority) (2011), 'Sektoruntersuchung Kraftstoffe' ('Sector Inquiry Fuels').

Byrne, D.P. and de Roos, N. (2014), 'Learning to Co-ordinate: A Study in Retail Gasoline', working paper.

Byrne, D.P. and G. Leslie (2013), 'Gasoline Pricing in the Country and the City', working paper.

Byrne, D.P., G. Leslie and R. Ware (2015), 'How do Consumers Respond to Gasoline Price Cycles?', *Energy Journal*, 36 (1).

Castanias, R. and H. Johnson (1993), 'Gas Wars: Retail Gasoline Price Fluctuations', *Review of Economics and Statistics*, 75 (1), 171-174.

De Roos, N. and H. Katayama (2013), 'Gasoline Price Cycles Under Discrete Time Pricing', *Economic Record*, 89 (285), 175-193.

Deltas, G. (2008), 'Retail Gasoline Price Dynamics and Local Market Power', *Journal of Industrial Economics*, 56 (3), 613-629.

Doyle, J., E. Muehlegger and K. Samphantharak (2010), 'Edgeworth Cycles Revisited', *Energy Economics*, 32 (3), 651-660.

Eckert, A. (2002), 'Retail Price Cycles and Response Asymmetry', *Canadian Journal of Economics*, 35 (1), 52-77.

Eckert, A. (2003), 'Retail Price Cycles and the Presence of Small Firms', *International Journal of Industrial Organization*, 21 (3), 151-170.

Eckert, A. and D.S. West (2004), 'Retail Gasoline Price Cycles across Spatially Dispersed Gasoline Stations', *Journal of Law and Economics*, 47 (1), 245-273.

Edelman, B. and M. Ostrovsky (2007), 'Strategic Bidder Behavior in Sponsored Search Auctions', *Decision Support Systems*, 43 (1), 192-198.

Edgeworth (1925), 'The Pure Theory of Monopoly', in Papers Relating to Political Economy, Vol. 1. London: MacMillan, 111-142.

Engle, R.F. and C.W.J. Granger (1987), 'Co-Integration and Error Correction: Representation, Estimation, and Testing', *Econometrica*, 55 (2), 251-276.

Federal Trade Commission (2011), 'Gasoline Price Changes and the Petroleum Industry: An Update'. Report by the Federal Trade Commission Bureau of Economics, September 2011.

Foros, O. and F. Steen (2008), 'Gasoline Prices Jump up on Mondays: An Outcome of Aggressive Competition?', working paper.

Foros, O. and F. Steen (2013), 'Vertical Control and Price Cycles in Gasoline Retailing', *Scandinavian Journal of Economics*, 115 (3), 640-661.

Harding, D. (2008). 'FoolWatch: A Case Study of Econometric Analysis and Evidenced Based Policy Making in the Australian Government', University of Munich MPRA Paper 16041.

Konkurransetilsynet (Norwegian Competition Authority) (2010). 'Det Norske Drivstoffmarkedet' ('The Norwegian Fuel Market').

Lewis, M. (2009), 'Temporary Wholesale Gasoline Price Spikes Have Long-Lasting Retail Effects: The Aftermath of Hurricane Rita', *Journal of Law and Economics*, 52 (3), 581-605.

Lewis, M. (2012), 'Price Leadership and Coordination in Retail Gasoline Markets with Price Cycles', *International Journal of Industrial Organization*, 30 (4), 342-351.

Lewis and Marvel (2011), 'When Do Consumers Search?', *Journal of Industrial Economics*, 59 (3), 457-483.

Lewis, M. and M. Noel (2011), 'The Speed of Gasoline Price Response in Markets with and without Edgeworth Cycles', *Review of Economics and Statistics*, 93 (2), 672-682.

Maskin, E. and J. Tirole (1988), 'A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles', *Econometrica*, 56 (3), 571-599.

Noel, M.D. (2002), 'Edgeworth Price Cycles in Retail Gasoline Markets', MIT Dissertation.

Noel, M.D. (2007a), 'Edgeworth Price Cycles, Cost-Based Pricing, and Sticky Pricing in Retail Gasoline Markets', *Review of Economics and Statistics*, 89 (2), 324-334.

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

Noel, M.D. (2008), 'Edgeworth Price Cycles and Focal Prices: Computational Dynamic Markov Equilibria', *Journal of Economics and Management Strategy*, 17 (2), 345-377.

Noel, M.D. (2009), 'Do Retail Gasoline Prices Respond Asymmetrically to Cost Shocks? The Influence of Edgeworth Cycles', *RAND Journal of Economics*, 40 (3), 582-595.

Noel, M.D. (2012), 'Edgeworth Price Cycles and Intertemporal Price Discrimination', *Energy Economics*, 34 (4), 942-954.

Noel, M.D. and L. Chu (2015), 'Forecasting Gasoline Prices in Markets with Edgeworth Price Cycles', *Energy Economics,* 51, 204-214.

Phillips, P.C.B. and S. Ouliaris (1990), 'Asymptotic Properties of Residual Based Tests for Cointegration', *Econometrica*, 58 (1), 165-193.

Stock, J.H., and M. W. Watson, (1993), 'A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems', Econometrica 61 (4), 783-820.

Wang, Z. (2009a), 'Station Level Gasoline Demand in an Australian Market with Regular Price Cycles', *Australian Journal of Agricultural and Resource Economics*, 53 (4), 467-483.

Wang, Z. (2009b), '(Mixed) Strategy in Oligopoly Pricing: Evidence from Gasoline Price Cycles Before and Under a Timing Regulation', *Journal of Political Economy*, 117 (6), 987-1030.

Zhang, X. (2006), 'Finding Edgeworth Cycles in Online Advertising Auctions', working paper.

Zimmerman, P.R., J.M. Yun and C. T. Taylor (2013), 'Edgeworth Price Cycles in Gasoline: Evidence from the U.S.', *Review of Industrial Organization*, 42 (3), 297-320.

Figure 1. Weekly Edgeworth price cycle in Windsor and daily Edgeworth price cycle in Toronto, 2007

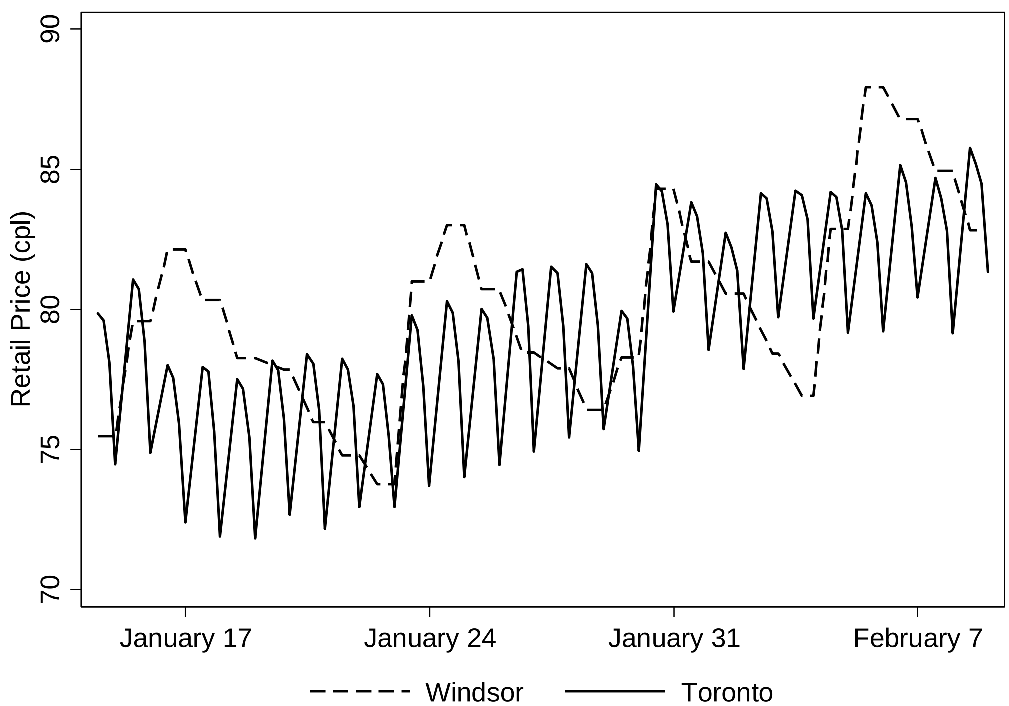

Figure 2. A theoretical Edgeworth price cycle with marginal cost equal to zero

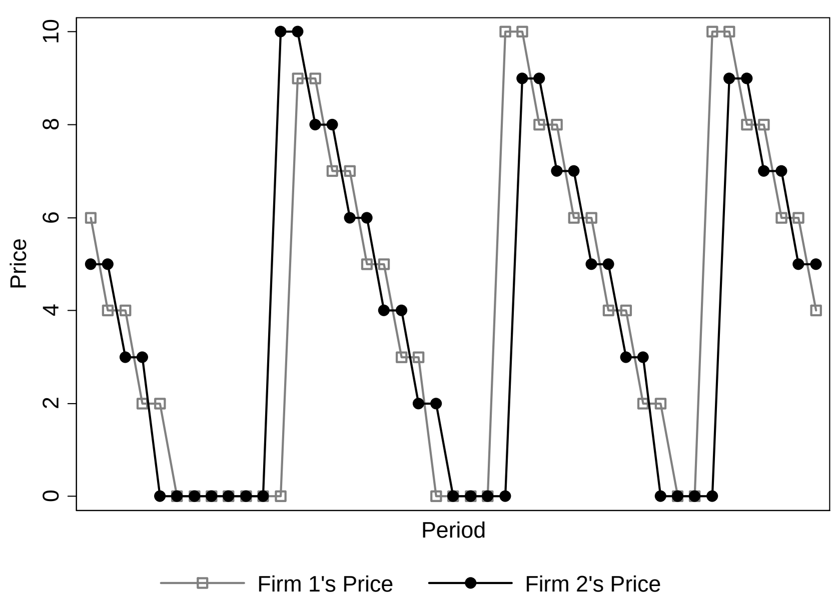

Figure 3. Daily Edgeworth price cycles cease in treatment cities after February 15th, 2007

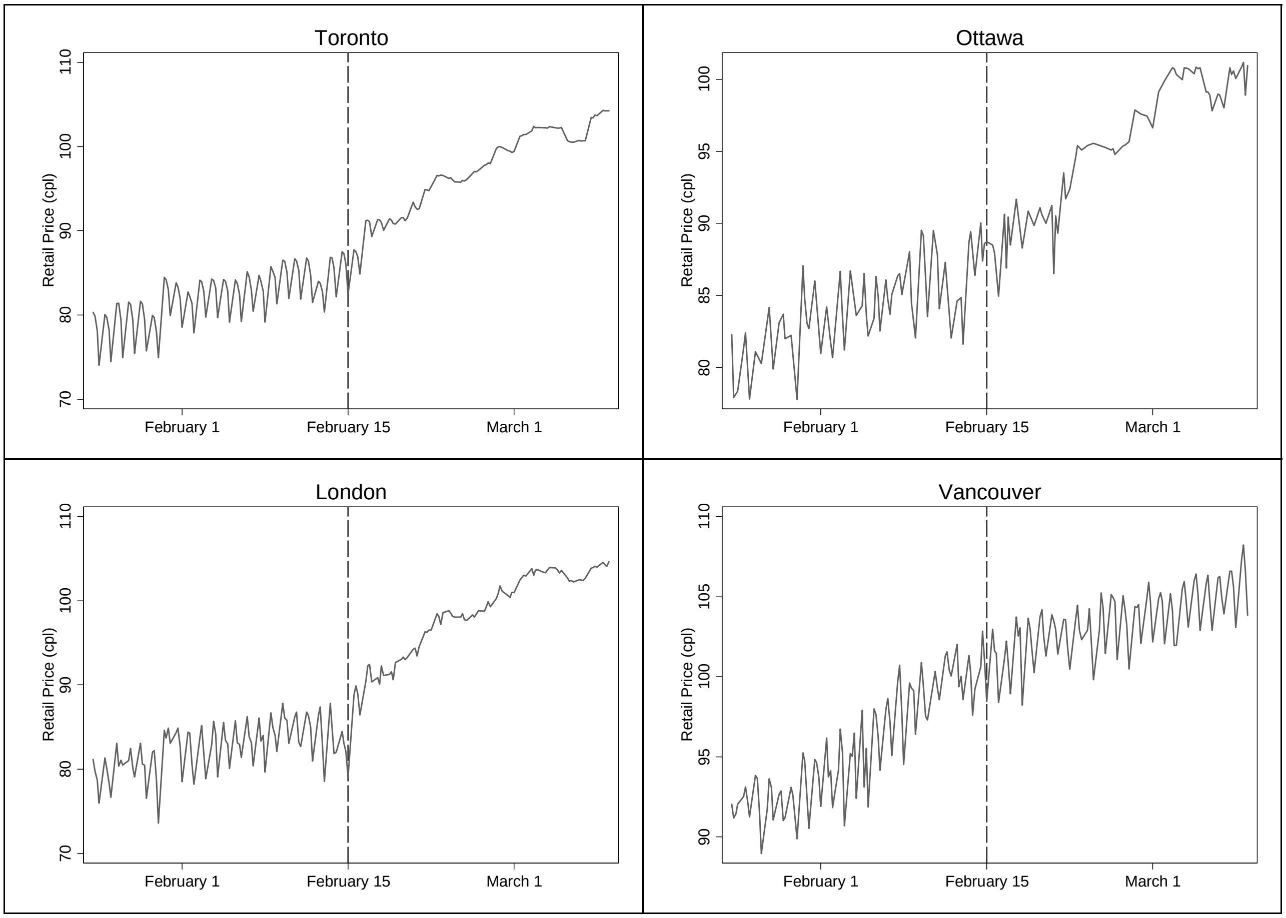

Figure 4. Margins Before and After the February 15, 2007 in Four Cities

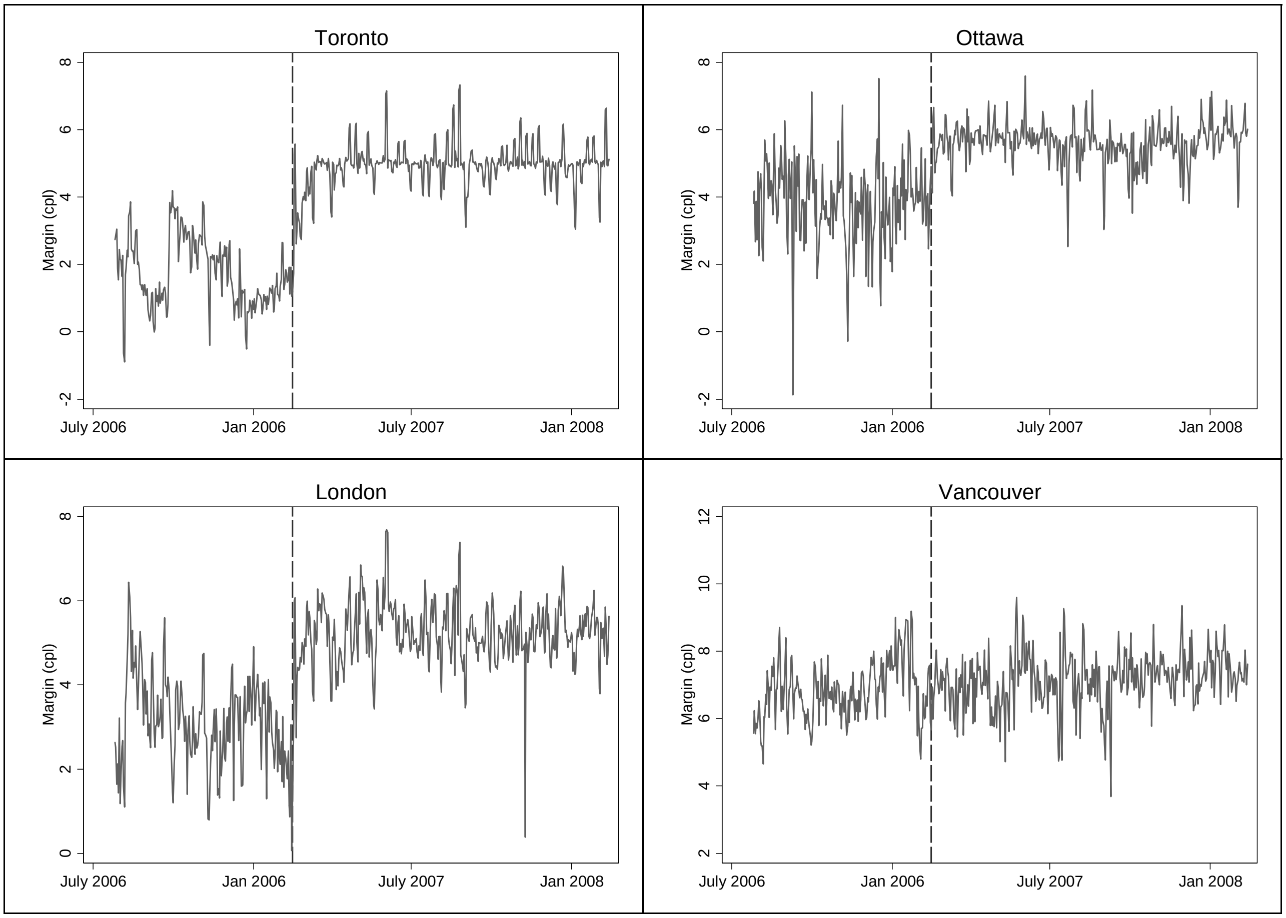

Table 1. Summary Statistics

	Mean	Std. Dev.	Minimum	Maximum
PRICE (Retail price per liter)	99.43	9.82	73.76	127.96
RACK (Rack price per liter)	62.58	7.77	43.80	85.93
TAX (Taxes per liter)	31.26	4.43	23.38	41.53
MARGIN (ex-tax)	5.61	3.20	-8.17	16.60

Retail and rack prices, taxes, and margins in Canadian cents per liter. (Approximate exchange rate over the sample: 1 Canadian dollar = 0.9164 US dollars).

Table 2. The Impact of Cycle Cessation on Price and Margins

Dep. Var. = PRICE	(1)	(2)	(3)
Post-Shock Period Begins:	March 26	April 30	June 15
RACK	0.988***	0.986***	0.991***
	(0.019)	(0.020)	(0.019)
ADJUSTMENT_PERIOD	1.010*	0.759	1.150***
	(0.563)	(0.383)	(0.237)
POST_SHOCK	1.258***	1.138***	1.271**
	(0.314)	(0.392)	(0.408)
TREATMENT*ADJUSTMENT_PERIOD	1.900**	2.477***	1.977***
	(0.703)	(0.543)	(0.372)
TREATMENT*POST_SHOCK	1.312***	1.090***	1.064***
	(0.274)	(0.317)	(0.312)
CITY AND MONTHLY DUMMIES	Y	Y	Y
R-SQUARED	0.99	0.99	0.99
NUM. OBS.	6064	6064	6064

Dep. Var. = MARGIN	(4)	(5)	(6)
Post-Shock Period Begins:	March 26	April 30	June 15
ADJUSTMENT_PERIOD	1.050	0.719	1.117**
	(0.586)	(0.398)	(0.236)
POST_SHOCK	1.177***	1.281***	1.203***
	(0.225)	(0.282)	(0.302)
TREATMENT*ADJUSTMENT_PERIOD	1.867**	2.245***	1.978***
	(0.664)	(0.524)	(0.372)
TREATMENT*POST_SHOCK	1.315***	1.092***	1.061***
	(0.271)	(0.311)	(0.315)
CITY AND MONTHLY DUMMIES	Y	Y	Y
R-SQUARED	0.88	0.88	0.88
NUM. OBS.	6064	6064	6064

Robust standard errors clustered by city in parentheses. *** Significant at 1% level, ** Significant at 5% level, * Significant at 10% level. City dummies included in all specifications obviating the need for "main" treatment effects.

Table 3. Robustness Specifications

*Cells show TREATMENT*POST_SHOCK coefficients from row specification*

Dep. Var. = PRICE

Post-Shock Period Begins:	March 26	April 30	June 15
Using Only Western Cities as Controls	1.360***	1.157**	1.101***
	(0.213)	(0.299)	(0.331)
Using Only Eastern Cities as Controls	1.389***	1.133	1.156*
	(0.554)	(0.622)	(0.564)
Using Only Quebec and Western Controls	1.391***	1.150**	1.114**
	(0.316)	(0.380)	(0.368)
Toronto Only vs. Controls	1.603***	1.390***	1.392***
	(0.235)	(0.288)	(0.276)
Ottawa & London Only vs. Controls	1.150***	0.928**	0.890**
	(0.242)	(0.288)	(0.277)
Reduced One-Year Event Window	1.790**	1.419*	1.491*
	(0.653)	(0.706)	(0.715)

Dep. Var. = MARGIN

Post-Shock Period Begins:	March 26	April 30	June 15
Using Only Western Cities as Controls	1.338***	1.137***	1.062**
	(0.195)	(0.268)	(0.325)
Using Only Eastern Cities as Controls	1.373***	1.123	1.141*
	(0.532)	(0.602)	(0.545)
Using Only Quebec and Western Controls	1.397***	1.154**	1.109**
	(0.309)	(0.369)	(0.370)
Toronto Only vs. Controls	1.605***	1.392***	1.391***
	(0.230)	(0.280)	(0.279)
Ottawa & London Only vs. Controls	1.153***	0.930**	0.885**
	(0.237)	(0.279)	(0.278)
Reduced One-Year Event Window	1.833**	1.483*	1.736**
	(0.646)	(0.692)	(0.706)

Robust standard errors clustered by city in parentheses. *** Significant at 1% level, ** Significant at 5% level, * Significant at 10% level. City dummies included in all specifications.

Table 4. Additional Robustness Specifications

*Cells show TREATMENT*POST_SHOCK coefficients from row specification*

Dep. Var. = PRICE

Post-Shock Period Begins:	March 26	April 30	June 15
Dynamic OLS Model	1.496***	1.306***	1.260***
	(0.214)	(0.217)	(0.216)
Error Correction Model	1.288***	1.076***	1.077***
	(0.356)	(0.359)	(0.361)
Include Cycle Position Measure	1.251***	1.029***	1.017***
	(0.256)	(0.304)	(0.300)
Collapse Data to Monthly Level	1.324***	1.125***	1.107***
	(0.269)	(0.292)	(0.286)
Collapse Data to Pre/Post Level	1.593***	1.292**	1.148**
	(0.390)	(0.423)	(0.440)
Regression Discontinuity	2.729***	2.617***	2.426***
	(0.100)	(0.169)	(0.247)

Dep. Var. = MARGIN

Post-Shock Period Begins:	March 26	April 30	June 15
Dynamic OLS model	1.497***	1.309***	1.285***
	(0.214)	(0.218)	(0.226)
Error Correction Model	1.332***	1.089***	1.126***
	(0.363)	(0.369)	(0.366)
Include Cycle Position Measure	1.255***	1.033***	1.014***
	(0.252)	(0.295)	(0.300)
Collapse Data to Monthly Level	1.318***	1.124***	1.122***
	(0.270)	(0.297)	(0.295)
Collapse Data to Pre/Post Level	1.389**	1.084*	1.107*
	(0.445)	(0.552)	(0.539)
Regression Discontinuity	2.571***	2.460**	2.303**
	(0.194)	(0.248)	(0.299)

Robust standard errors clustered by city in parentheses. *** Significant at 1% level, ** Significant at 5% level, * Significant at 10% level. City dummies included in all specifications.