Calendar Synchronization of Gasoline Price Increases

Michael D. Noel*

May 18, 2018

Abstract

In many retail gasoline markets with Edgeworth price cycles, large and regular price increases occur on the same day of the week every week, i.e. they are calendar synchronized. In this article, I test whether calendar synchronization leads to higher or lower consumer expenditures on gasoline compared to a world with cycles but without calendar synchronization. On one hand, firms may attempt to trigger price increases just prior to periods of normally high demand. On the other, consumers may be better able to predict and shift purchases to low price days of the cycle. Using high-frequency gasoline volume data and matching it to high-frequency price data, I find that the latter effect dominates. All else equal, consumer expenditures on gasoline fall with calendar synchronization in the study markets. I also calculate intertemporal price elasticities and find them to be high.

JEL Classification Codes: L1, L2, L4, M3, Q4

Keywords: Edgeworth price cycles, retail gasoline, dynamic oligopoly, calendar effects, prediction, quantity-weighting

1 Introduction

Economists have long been interested in the striking pattern found in retail gasoline prices in many cities around the world. In these cities, retail gasoline prices follow a high-frequency, asymmetric price cycle in which prices rise quickly, often 10% or more over the course of a day or two, and then fall back more slowly, often over a week or two. The cycle repeats itself again and again even in the absence of wholesale price changes. In some cities, the asymmetric cycle can be especially rapid, with prices rising and falling asymmetrically in a matter of only hours instead of days or weeks.

Figure 1 shows an example of the retail gasoline price cycles in Perth, Australia, in 2012 and 2013. Cycles similar to these have been documented in other Australian cities, Canada, the United

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

States, and various European countries. The leading explanation for the cycles is that they are Edgeworth price cycles, the outcome of a dynamic oligopoly game formalized theoretically by Maskin and Tirole (1988). Empirical studies finding support for the Edgeworth price cycle theory include Eckert (2003), Noel (2007a), Noel (2007b), Atkinson (2009), Lewis (2009), Lewis (2012), Wang (2009a), Wang (2009b), Doyle et al. (2010), and others.

While the potential causes of the cycles have been well studied, little attention has been given to the precise timing of the ups and downs of the cycles. In particular, the timing of the relenting phase (i.e. when prices increase) has been a controversial topic among consumers, regulators, and economists. A close inspection of Figure 1 reveals the issue. The cycles in Perth, as in many other cities, are not just *about* a week long – they are *exactly* a week long. Peaks and troughs are synchronized to the calendar week. Troughs always occurred on Wednesdays and peaks always occurred on Thursdays during this period. Weekly synchronized cycles such as these have been observed in other major Australian cities as well (Noel and Chu (2015), Valadkhani (2013)), along with Norway (Foros and Steen (2013)), the United States (Lewis (2012)) and Canada (Noel (2007a), Byrne et al. (2015)). More recently, daily-frequency price cycles synchronized to the hour of the day have been observed in Canada (Atkinson et al. (2014), Noel (2015)) and Germany (Haucap et al. (2016), Siekmann (2017)).

This invites the obvious question – conditional on the presence of retail gasoline price cycles, does the calendar synchronization feature of those cycles benefit or harm consumers, all else equal? Are firms timing the relenting phases so that the "high price days" now coincide with the days of the week when gasoline demand would normally be high anyway? Or are consumers using the better predictability of the cycle troughs to shift their purchases to "low price days"? The first potential effect is the "relenting-phase-shifting" effect by firms, which contributes to higher consumer expenditures, and the second is the "intertemporal switching effect" by consumers, which contributes to lower consumer expenditures. In this article, I use high-frequency data on gasoline prices and gasoline volumes to estimate the net effect of calendar synchronization on the prices consumers pay, conditional on the presence of price cycles, and holding all else about the cycles and the marketplace equal.

The question I am asking is different than the question of whether price cycles are beneficial or harmful to consumers as a whole, compared to a world with no cycles. Recent research suggests that cycles are in fact beneficial overall (Noel (2002), Doyle et al. (2010), Lewis and Noel (2012), Zimmerman et al. (2013), Noel (2015), Siekmann (2017)). These articles do not consider the calendar synchronization feature itself, however, and that is my purpose here. To my knowledge, this is the first study to do so. I envision two otherwise identical worlds – one in which retail gasoline price cycles are present and synchronized to the calendar, and a counterfactual world in which the same basic cycles are present but not synchronized to the calendar.

Past studies have made conjectures about how calendar synchronization might affect consumer expenditures. Noel (2007a) examines Canadian cycles and suggests that firms may have an incentive to trigger relenting phases at certain times of the week, prior to days with normally high demand. Foros and Steen (2013) examine cycles in Norway and conjecture that firms may be strategically coordinating relenting phases on specific days of the week to increase their profits. In contrast, Noel (2012) and Noel and Chu (2015) conjecture that calendar synchronization might be beneficial to consumers by making it easier for them to predict cycle troughs and buy at relatively low prices. In policy discussions, a sort of folk theory has developed in recent years that begins with the presumption that calendar synchronization is an undesirable feature of price cycles that is necessarily harmful to consumers.

In spite of all the conjectures and presumptions, no study has been able to actually test them. One obstacle has been the difficulty in obtaining high-frequency gasoline volume data (i.e. gasoline quantity data). Publicly available volume data is rarely available at better than a monthly frequency, but to address the question at hand in a world with weekly price cycles, volume data at the daily frequency level is needed to match daily changes in volumes with daily changes in prices.[1]

I overcome this difficulty with access to a limited amount of high-frequency daily gasoline volume data which I can then match to daily gasoline price data – for all five major Australian capital cities of Adelaide, Brisbane, Melbourne, Sydney, and Perth, over several years. High-frequency volume data is largely missing in the gasoline literature and this article represents a rare exception in that regard.

I offer several contributions. The first is the application methodology itself. In place of pre-

[1] In a world with daily retail price cycles, even more frequent volume data is necessary, e.g. hourly.

sumptions and conjectures, this article lays out a simple and transparent method for evaluating the effects of calendar synchronization and volume shifting for the first time. It compares consumer expenditures in the presence of calendar synchronized price cycles to consumer expenditures still in the presence of price cycles but absent the calendar synchronization feature (holding all else equal). The key simplifying insight is that expected prices in a world without calendar-synchronization are day-of-the-week invariant and, as a result, the relevant comparison boils down to a deceptively simple comparison of quantity-weighted and quantity-unweighted average prices. The simplicity of the methodology is its virtue. It can be easily implemented by competition authorities and others evaluating the issue of calendar synchronization and volume shifting in their own jurisdictions.

The second contribution lies in the main results. The results ultimately show that the intertemporal switching effect by consumers dominates any relenting-phase-shifting effect by firms meaning that calendar synchronization, holding all else equal, reduces consumer expenditures. The results are consistent and statistically significant across all five major Australian cities. While the results may or may not be omnipresent or as strong in other jurisdictions, the analysis makes it clear that past presumptions surrounding presumed negative effects of calendar synchronization need to be rethought. In fact, given the strength and consistency of the results across all five major cities, the results tend to support the opposite prior. Calendar synchronization may be beneficial to consumers after all, in expectation.

The third contribution of this article lies in the high-frequency gasoline volume data itself. Such data tends to be proprietary, difficult to access, and is rarely seen in the literature. Yet high-frequency gasoline volume data is important for understanding what consumers really pay in a world where retail gasoline price cycles are present and where prices frequently change. In my main results, I find that the reduction in consumer expenditures with calendar synchronization is due to a substantial degree of volume shifting by consumers to periods of low prices, which can only be observed with high-frequency volume data. An important implication of this result is that the unweighted average gasoline price, which is so often used in making price comparisons in the gasoline literature, can significantly overstate the prices that consumers actually pay when there are retail price cycles. In this article, I emphasize the importance of incorporating high-frequency volume data into the analysis whenever possible, while at the same time recognizing the difficult challenge that it is.

Finally, a fourth contribution is in the calculation of gasoline price elasticities, in particular, intertemporal price elasticities of demand in response to a *temporary* and recurring change in price. These are different than the short and long run gasoline price elasticities in response to a *permanent* change in price that have been well studied in the literature. Intertemporal elasticity estimates are rarely encountered. The reason is that the calculation of intertemporal elasticities require large, predictable, and temporary "sales" on gasoline which are uncommon outside the context of retail gasoline price cycles. In the context of gasoline price cycles, such a calculation – in this case, measuring consumers' willingness to shift purchases to different days of the week along the cycle – is possible but requires high-frequency volume data which, again, is difficult to obtain. In this study, having obtained the necessary data, I find that intertemporal price elasticities are very high and conclude that consumers are very responsive to temporary and predictable changes in gasoline prices. In large numbers they shift their purchases across the days of the week to take advantage of the known temporary "sales" at the troughs of the cycles.

2 Literature

Retail gasoline price cycles are often referred to as "Edgeworth price cycles". Edgeworth price cycles were first postulated by Edgeworth (1925), and formalized theoretically by Maskin & Tirole (1988). After the discovery of retail gasoline price cycles in several Canadian cities, Eckert (2003) and Noel (2008) extended the Edgeworth price cycle model to a variety of more complex situations, with the realities of retail gasoline markets in mind.

In the standard model, two infinitely-lived profit-maximizing firms compete in a homogeneous Bertrand pricing game by setting prices (from a discrete price grid) in an alternating fashion. The strategies are allowed to depend only on the payoff-relevant state, i.e. current cost and demand and the price set by the other firm in the previous period. The Markov Perfect equilibrium strategies are given by $(p_t^1)^* = R^1(p_{t-1}^2)$, $(p_t^2)^* = R^2(p_{t-1}^1)$ where p_{t-1}^j is the price chosen by firm j in period $t-1$ which remains in effect in period t.

Let $V^1(p_{t-1}^2)$ be Firm 1's value function when Firm 2 adjusted its price to p_{t-1}^2 in the previous period, and Firm 1 adjusts its price in the current period. Let $W^1(p^1_{s-1})$ 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 δ is the discount factor. Similar equations exist for V^2 and W^2.

Maskin and Tirole (1988) show that two sets of equilibria are possible – focal price equilibria and Edgeworth price cycle equilibria. Equilibrium strategies that result in Edgeworth price cycles take the form:

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

Along the path of cycle, the two firms undercut one another until the price gets close to marginal cost, at which point one firm drops its price immediately to marginal cost and the other follows. Firms then play a war of attrition, in which each mixes between raising its price higher or keeping its price at marginal cost for another round. Eventually one firm raises its price, the other follows, and undercutting resumes from the top of the cycle almost immediately. The cycle repeats over and over in this fashion.

Eckert (2003) extends the model to account for unequal sharing rules at equal prices and Noel (2008) extends the model to a variety of different situations, including varying elasticities of demand, marginal costs, degrees of differentiation, capacity constraints, firm asymmetries, and triopoly situations. The extensions generate a series of additional testable predictions about where the cycles are likely to occur and how they are likely to function over time.

Empirically, Allvine and Patterson (1974) were first to note the presence of retail gasoline price cycles in several U.S. cities in the 1960s and 1970s while Castanias and Johnson (1993) were first to note their similar appearance to Edgeworth price cycles. Since then, a large empirical literature has emerged to study retail gasoline price cycles in various markets and test the predictions of the Edgeworth price cycle model. Studies finding support for the model include Eckert (2003), Noel (2007a), Noel (2007b), Atkinson (2009), Lewis (2009), Lewis (2012), Wang (2009a), Wang(2009b), Doyle et al. (2010), and others.[2]

One aspect of Edgeworth price cycles that is especially relevant to this study is the frequent appearance on the equilibrium path of what Noel (2008) calls "Delayed Starts" and "False Starts". Noel shows that Delayed Starts and False Starts are normal parts of most cycling equilibria and become more frequent when there is greater market uncertainty. A Delayed Start occurs when a firm (or set of firms) increases price to the top of the cycle, but the remaining firms do not follow the increase very quickly. The first firm is temporarily stranded with a high price at the top of the cycle, making few sales, and incurring losses until other firms raise their prices as well (if they do at all). Noel shows that, if a Delayed Start occurs and goes on for too long, the best course of action for the first firm is to abandon its price increase altogether and rejoin the war of attrition at the bottom of the cycle, causing what is known as a False Start. Atkinson (2009) first verified the existence of False Starts empirically and, since then, high-frequency price gasoline price data reveals that False Starts are commonplace in many markets (e.g. Noel (2015), Byrne and de Roos (2017)).

Relenting phases in general, and Delayed and False Starts in particular, are risky and costly to the firms that raise prices first. It is never certain that a relenting phase, once started, will complete quickly, or that it will complete at all. One risk to the firm that raises price first is that late movers

[2] One alternate theory is that prices may simply be rising and falling to reflect the deterministic pattern of demand over the course of a week (or the course of a day). The explanation is easily debunked, however, as discussed by Noel (2007b). In short, the normal pattern of weekly demand across most major urban centers follows an inverted U-shaped pattern with highest demand during the workweek and lowest demand on weekends. It is not the case – and is not plausible – that demand would suddenly rise by a large amount on one particular day of the week and then fall by small and consistent amounts every other day of the week, as required by this story. There are also no intraday "bumps" along weekly cycles to account for rush hour demand spikes. Moreover, the specific day on which relenting phases occur is different in different cities and can change over time even within the same city, inconsistent a demand story. In some cities (e.g. Canada), relenting phases occur on different days each week meaning that a given day of the week can oscillate between peak and trough days, also inconsistent with a demand story. Other cities have two relenting phases in a week (Mondays and Thursdays in Norway) and some cities with daily cycles have relenting phases in the morning (Canada) while others have relenting phases in the evening (Germany), all inconsistent with double-peaked weekday demand and single-peaked weekly demand. In short, the cycles are not caused by determinstic weekly and daily demand patterns. The question of interest in this study is whether deterministic demand patterns can influence the exact timing of relenting phases.

may be slow in learning and reacting to the initial price increases taking place somewhere in the city. Another is that late movers may themselves become engaged in a competitive battle over who raises prices earlier and who delays longer, all the while the first firm continues to lose sales at the top of the cycle. Another risk is that wholesale prices may fall in the middle of a relenting phase causing late movers to not follow the price increase. Noel (2008) shows that unexpected decreases in wholesale prices in the midst of a marketwide relenting phase are a common cause of Delayed and False Starts. As falling wholesale prices create more space for late movers to continue undercutting one another at the bottom of the cycle, the firm that raised prices first must either wait an extended time at too high a price or cut its losses and abandon its price increases in a False Start. Regardless of the cause, the first firm runs the risk of developing an unwanted reputation as a high priced firm, a risk that is exacerbated when relenting phases are longer and Delayed and False Starts are more common.

The risks translate to lower market shares and profits for the firm that raises prices first. Noel (2008) shows that firms that raise prices first earn substantially lower profits than firms who raise prices later, and that the profit differential grows larger when relenting phases are more drawn out and more prone to Delayed Starts and False Starts.

Faced with these risks, it is reasonable for a firm that increases prices first to try to minimize the length of the relenting phases and the frequency of False Starts. One way to do this is to try to start relenting phases at a similar time of day or a similar day of the week – i.e. to calendar synchronize the cycles. Doing so would potentially make the firm's price increases more predictable and visible to other firms, reduce the length of the relenting phases, and make it easier for other firms to gauge when a False Start may be imminent, reducing the frequency of these costly events.

So even if calendar synchronization were beneficial to consumers as a whole (by making it easier for consumers to predict periods of low prices), and even if it were harmful to firms as a whole, it remains a reasonable course of action for the firm that raises prices first to try to calendar synchronize the cycles and steal some market share and profits back from the late movers. At the same time, it remains a reasonable course of action for the late movers to continue following the first firm in relatively short order to avoid a False Start, since late movers still benefit from winning the war of attrition (even if not by as much), and a False Start would simply reopen that war of attrition again.

To be clear, cycles that are calendar synchronized does not mean cycles that are uncompetitive. With or without calendar synchronization, relenting phases continue to occur only when they are competitively necessary. If calendar synchronization causes the cycle period to change, the other characteristics of the cycle, amplitude in particular, competitively compensates in the expected way. As an example, note how in Figure 1 the price increases in the relenting phase are larger when wholesale prices are increasing than when they are decreasing. When wholesale prices are increasing and are expected to increase, a larger price increase is needed for the normal rate of undercutting to reach the "bottom" of the cycle one week later. When wholesale prices are decreasing, a smaller price increase is needed for the normal rate of undercutting to reach the bottom of the cycle one week later. The size of the initial price increase (which determines the amplitude of the next cycle) thus competitively adjusts so that the "bottom of the cycle" will be reached in expectation on the usual day. This is a general result with cycles – with a more stable period, amplitudes competitively adjust to compensate and with a more stable amplitude, periods competitively adjust to compensate. In this way, cycle periods and amplitudes are interconnected, and relenting phases continue to be competitively necessary even though the cycles are synchronized to the calendar. The main advantage of a firm trying to synchronize cycles to the calendar is simply that it limits its market share and profit losses to the late movers.

Calendar synchronization turns out to be common in many countries that experience retail gasoline price cycles (Noel (2007b), Lewis (2012), Foros and Steen (2013), Noel and Chu (2015), Valadkhani (2013), Byrne and de Roos (2015), and Siekmann (2017)). While it may seem surprising that little research has been done to evaluate it, a main roadblock has been the lack of availability of high-frequency gasoline volume data. For the first time with this study, with high-frequency volume data in hand, I move beyond the conjectures and presumptions and explicitly test for the effects of calendar synchronization on consumer expenditures in these markets.

3 Data

I examine the effects of calendar synchronization using price and volume data for five major Australian cities – Adelaide, Brisbane, Melbourne, Sydney and Perth – from July 1, 2009 to June 30, 2010 and from July 1, 2012 to June 30, 2013. The former period I refer to as the 2009-2010 period and the latter I refer to as the 2012-2013 period. The first four cities, when referencing them collectively, I call the "Eastern Cities". As shown below, all five cities experienced calendar synchronized price cycles throughout the 2009-2010 period. While all five continued to have price cycles throughout the 2012-2013 period, only the cycles in Perth continued to be calendar synchronized.

I first collect retail and wholesale gasoline prices at a daily level, for each city and each of the two year-long time periods. The retail price data is based on the average daily price of regular grade gasoline in each city and is obtained from the Australian Competition and Consumer Commission ("ACCC"). The wholesale price data is based on the daily Terminal Gate Price ("TGP") in each city, also known as the rack price, and is obtained from Orima Research. Both are measured in Australian cents per liter and adjusted to 2012Q2 dollars, with taxes removed.

Unique to this article is the use of a limited amount of high-frequency daily gasoline volume data that can be matched to the price data. High-frequency volume data is difficult to obtain and rarely seen in the literature.[3] The volume data was obtained from the ACCC, which in turn obtained them directly from major firms as part of its mandated annual petroleum industry review. Firms do not otherwise provide this information, and the raw data is not released by the ACCC. Fortunately, since the cycles studied here are weekly cycles and since they are synchronized to the day of the week, week after week, the key volume data needed for this study is a measure of gasoline volumes in each city for each different day of the week in each study period. This can then be matched with daily retail prices that themselves vary substantially by the day of the week, but in a consistent way, week after week and cycle after cycle (controlling for costs). The relevant variation for my comparisons – i.e. the within-cycle variation in prices and the within-cycle variation in volumes along the path of the cycle – is preserved and matched. I obtain from the ACCC the set of point

[3] Exceptions include Wang (2009a) and Barron, Umbeck and Waddell (2008), who were given access to volume data for a small set of stations cooperating with the study.

estimates from a fixed effect regression of daily gasoline volumes on a complete set of day of the week indicator variables, for each city and for each separate time period.

I report summary statistics for prices and volumes in Table 1.

4 Methodology

The primary goal is to determine whether consumer expenditures for a given volume of gasoline are higher or lower when retail gasoline price cycles are synchronized to the day of the week compared to an identical world with retail gasoline price cycles that are not synchronized to the day of the week. As shown below, the price cycles in all five Australian cities in 2009-2010 and in Perth in 2012-2013 are examples of cycles that are synchronized to the day of the week. The troughs generally occurred on Wednesdays and the peak occurred a day or two later, depending on the city. There were clear and predictable high price days and low price days of the week in each city.

There are two potential effects. First, consumers have the ability to shift purchases to the easily predictable "low price days" of the week closest to the trough. This "intertemporal switching effect" causes volumes to increase on low price days and contributes to lower consumer expenditures. Second, firms could potentially attempt to move relenting phases just prior to what would otherwise be normally high demand days. This "relenting-phase-shifting effect" would cause the "high price days" of the cycle to coincide with days when volumes would normally be high and contribute to higher consumer expenditures. Whether consumer expenditures increase or decrease overall, all else equal and compared to the relevant counterfactual world, depends on the relative strength of the two effects.

The counterfactual world is a world in which there are still price cycles but that are not synchronized to the day of the week. In this world, relenting phases do not occur on the same day each week and are random in the sense that the expected price for any one day of the week is the same as any other day of the week. That is, there are no high price days or low price days. Firms have not attempted or were not able to synchronize relenting phases to any given day of the week and consumers cannot predict high price days and low price days because these do not exist. In other words, in the counterfactual world, prices are uncorrelated with the day of the week and price is day-of-the-week invariant.

I wish to compare average weekly consumer expenditures with cycles in the presence of calendar synchronization to average weekly consumer expenditures with the same cycles but in the absence of calendar synchronization. To hold all else equal, I assume that the price distributions in the calendar synchronized and unsynchronized worlds share the same overall mean, but the price distributions themselves can differ and the distribution of gasoline volumes can also differ. I normalize volumes so as to compare the same amount of gasoline in each case, and the obvious normalization is to scale volumes to one liter. With this normalization, the comparison is mathematically equivalent to comparing the weekly quantity-weighted average price with calendar synchronization to the weekly quantity-weighted average price without calendar synchronization.

The reader will note that the counterfactual cycles, i.e. absent calendar synchronization, are not directly observable in the same cities at the same time that the actual calendar synchronized cycles are present. One might look to other cities or years for roughly comparable cycles to use as the counterfactual but, perhaps surprisingly, this is not necessary for the comparisons I wish to make. The weekly quantity-weighted average price in the unobserved counterfactual world is easily estimable with actual data by appealing to a key insight.

The insight is that the weekly quantity-weighted average price in the counterfactual world without synchronized price cycles is actually the same as the weekly unweighted average price. The reason is that, by definition, the expected price on any given day of the week is the same as on any other when cycles are not calendar synchronized. There are no high price days or low price days of the week, prices are uncorrelated with the day of the week, and price is day-of-the-week invariant. Therefore, unobserved quantities in the counterfactual world do not actually impact the calculation of quantity-weighted average prices in that world. Any set of quantity weights multiplied by the day-of-the-week-invariant average price returns the same day-of-the-week-invariant average price back as its result:

$$\sum_{i=1}^{7} a_i \overline{p} = \overline{p} \quad \forall \ a_i \tag{4}$$

for any set of quantity weights over the days of the week, a_i, and where $\overline{p}$ is the day-of-the-week invariant price. Not observing quantities in the counterfactual world is immaterial.

In contrast, in the actual world with calendar-synchronized prices, there are high price days and low price days and quantity weights do matter. Fortunately, quantities in the actual world are directly observed.

In light of this, the comparison of normalized consumer expenditures in a world with calendar-synchronized cycles to an identical world with the same cycles but without the calendar synchronization feature boils down to a deceptively simple calculation. It is equivalent to comparing weekly quantity-weighted average prices to weekly unweighted average prices for a given city and a given time period, holding all else equal but for the calendar synchronization feature. The test is given by:

$$\sum_{i=1}^{7} q_i p_i \gtrless \sum_{i=1}^{7} a_i p = \overline{p} \tag{5}$$

where q_i are the actual volume weights by day of the week with calendar synchronization and a_i are any set of weights. If the left hand side of the inequality is smaller than the right hand side, the weekly quantity-weighted average price in the actual world is lower than the weekly unweighted average price and consumers benefit from the calendar synchronization feature of the cycles, all else equal. The effect of intertemporal switching behavior by consumers dominates the effect of potential relenting-phase-shifting by firms and consumer expenditures fall. If instead the left hand side is larger, the weekly quantity-weighted average price is greater than the weekly unweighted average price, and consumers pay more with the calendar synchronization feature, all else equal. The effect of relenting-phase-shifting dominates the intertemporal switching effect, and consumer expenditures rise.

The approach has two important advantages. First, it is simple and transparent and can be easily used by competition authorities to evaluate calendar synchronization in their own jurisdictions. Second, it requires only data on prices and quantities for the relevant time period with calendar synchronization. It does not require finding examples of hopefully near identical markets and cycles without the calendar synchronization feature to compare them to. In many jurisdictions, the latter does not exist and, when they do, the analysis requires either long-panel or cross-sectional comparisons with a potentially difficult set of confounding omitted variables to control for. That is not to say that such an inquiry would not be useful (it can be used to relax the same-mean assumption, for example), but is more suited to a later stage of the investigation.

I conduct comparisons of quantity-weighted normalized consumer expenditures, with and without calendar synchronization – or equivalently, comparisons of weekly quantity-weighted average prices to weekly unweighted average prices – for all five major Australian cities of Adelaide, Brisbane, Melbourne, Perth and Sydney from July 2009 to June 2010 (the last year that cycles were synchronized to the day of the week in all the cities), and for Perth from July 2012 to June 2013 (which continued to experience calendar synchronized cycles).

As a preliminary exercise, I first need to establish that the retail gasoline price cycles in these cities were weekly and synchronized to the day-of-the-week. To do so, I estimate the core parameters of the cycle for each city and each relevant time period, including average price changes in the relenting and undercutting phases and the transition probability matrix between phases. I then derive the cycle characteristics for each city and time span, such as period, amplitude, and asymmetry, from the core parameters.

I define a relenting phase as a sequence of consecutive city-average price increases with a cumulative price increase of at least h^R cents per liter and the undercutting phase as a sequence of consecutive city-average price decreases with a cumulative price decrease of at least h^U.[4] I set $h^R = h^U = 3$.[5]

I define $\Delta RETAIL^i_{mt} = RETAIL_{mt} - RETAIL_{m,t-1}$ as the change in the retail price in city m at time t conditional on phase i and define the city-specific mean of the series as $\alpha^i_m = E(\Delta RETAIL^i_{mt})$.

For the transition probability matrix, I define I_{st} to be equal to R or U when the city is in a relenting or undercutting phase respectively, and λ^{ij}_{mt} to be the probability that market m transitions from phase i to phase j in t:

$$\lambda^{ij}_{mt} = \Pr(I_{mt} = j \mid I_{m,t-1} = i), \qquad i, j = \{R, U\} \tag{6}$$

[4] While each station tends to increase its price in a relenting phase all at once (e.g. Noel (2007b)), relenting phases defined over city-average prices can take multiple days to complete, due to averaging.

[5] The cycles are clearly delineated in the data and other reasonable values for h^i yield substantially similar results.

with λ_m^{ij} denoting the city-specific mean of the series.

Non-linear transformations of the core parameter vector $\{\alpha_m, \lambda_m^{ij}\}$ yield measures of the characteristics of the cycle. Following the calculations in Noel (2007a) and suppressing subscripts to simplify notation, I calculate the expected duration of phase i in each city as:

$$E(DURATION^i) = (1 - \lambda^{ii})^{-1} \tag{7}$$

and the expected period of the cycle as:

$$E(PERIOD) = (1 - \lambda^{RR})^{-1} + (1 - \lambda^{UU})^{-1} \tag{8}$$

I calculate cycle amplitude as:

$$E(AMPLITUDE) = \alpha^R (1 - \lambda^{RR})^{-1} \tag{9}$$

and cycle asymmetry as:

$$E(ASYMMETRY) = -\alpha^R / \alpha^U \tag{10}$$

Once a weekly cycle is established, it is straightforward to show that the cycles are synchronized to the calendar week using a regression of daily prices on day-of-the-week fixed effects.

5 Results

5.1 Weekly Cycle Characteristics

I now establish the existence of weekly price cycles in the five major Australian capital cities in 2009-2010 and again in Perth in 2012-2013. The core cycle parameters for each city during 2009-2010 are reported in Specifications (1) through (5) in Table 2. The average one-day price increase in a relenting phase ranged from 2.9 cents per liter in Perth to 7.2 cents per liter in Adelaide (α^R) and the average one-day price decrease in an undercutting phase ranged from 1.3 cents per liter in Perth to 2.7 cents per liter in Adelaide (α^U). The transition probability from a relenting phase to another relenting phase (λ^{RR}) was approximately 50% in the Eastern Cities and close to 60% for Perth (λ^{RR}). The transition probability from an undercutting phase to another undercutting phase (λ^{UU}) was approximately 80% in all cities, with relatively less variation.

For Perth in 2012-2013 (reported in Specification (6)), the average one day price increase was 8.9 cents per liter and the one-day price decrease was 1.5 cents per liter. The transition probability from a relenting phase to another relenting phase was insignificantly different from zero (i.e. two consecutive city-average price increases were rare in Perth at that time), and the transition probability from an undercutting phase to another undercutting phase was 83%.

Using the aforementioned transformations, Table 3 reports the derived characteristics of the cycles, for each city in 2009-2010 and for Perth in 2012-2013. In 2009-2010, relenting phases in the Eastern Cities were generally complete within two days (with coefficients ranging from 1.9 days to 2.0 days) while in Perth it sometimes took a third day, averaging 2.5 days overall.[6] The undercutting phase was generally complete in about five days (coefficients ranging from 5.1 to 5.5 days), yielding a cycle period insignificantly different from seven days in every city (coefficients ranging from 7.0 to 7.9).[7] Amplitudes varied, from a low of 7.0 cents per liter in Perth to a high of 13.6 cents per liter in Adelaide, and there was a statistically significant asymmetry in all cities.

In 2012-2013, relenting phases in Perth were generally complete in one day (coefficient=1.038) and undercutting phases in the remaining six days of the week (coefficient=5.981), resulting in a cycle period of almost exactly seven days (coefficient=7.019). The amplitude was 9.3 cents per liter and the asymmetry more than doubled from the earlier time frame.

To confirm that these cycles were not only weekly on average but that they were also synchronized to specific days of the week, Table 4 contains a series of regressions of prices on day of the week indicator variables, with Sunday being the omitted day. It shows that the cycles were indeed

[6] In Perth in late 2009, there was often a double trough on Tuesdays and Wednesdays, with the absolute trough sometimes occurring on either day. If occurring on a Tuesday, there was a very small price increase on Wednesday (typically less than a penny as a few stations began the relenting phase process late on Wednesday night). This results in a higher transition probability from a relenting phase to another relenting phase in Perth (λ^{RR}) while at the same time produces a smaller average daily price increase conditional on the relenting phase (α^{R}).

[7] The point estimates modestly exceed seven in most cases because of the presence of False Starts. From time to time, and often the result of falling wholesale prices, the relenting phase began but did not complete when it might normally have done so, causing prices to quickly return to the trough, and leading to an extended two-week cycle. There were no False Starts in Adelaide in 2009-2010 or in Perth in 2012-2013, so the point estimates in those cities are almost exactly seven. When h^R and h^U are decreased from three down towards zero, the point estimates on the cycle period approach seven in all cities, and the amplitude decreases to compensate in the expected way.

synchronized to the day-of-the-week in all cases. In all cities in 2009-2010, prices fell early in the week and were lowest on Wednesdays (coefficients range from -3.4 in Perth to -7.0 in Adelaide). Average prices began to rise on Thursdays, peaking on Fridays in Adelaide, Perth, and Sydney and on Saturdays in Brisbane and Melbourne, before starting to fall again. In Perth in 2012-2013, prices were at the trough on Wednesdays and at the peak on Thursdays.

The results show that there were clear low price days and high price days of the week and a strong correlation between expected prices and the day of the week. On one hand, consumers may use the predictability to shift purchases to low price days. (They would pay 7.0 to 13.6 cents per liter less if they bought on a Wednesday instead of a few days later.) On the other, firms may attempt to move relenting phases at certain times of the week to shift high price days to days of normally high demand.

To preview the effect that the calendar synchronized cycles had on the day-of-the-week pattern of volumes, Figure 2 plots volumes in each city over the relevant time periods and by day of the week. The figure shows that, in each city in each relevant time period, volumes peaked on Wednesdays, the same day as the trough. Wednesday volumes were approximately double what Friday volumes were, when prices were at or near the peak. While part of the inverted U-shaped pattern will reflect normal demand differences (it is known that weekend demand is lower that weekday demand, all else equal), the dramatic collapse in volumes from Wednesday to Thursday and Friday is suggestive that a substantial number of consumers may be effective in timing the cycle.

5.2 The Effects of Calendar Synchronization

The main results of this article are contained in Table 5. The table presents comparisons of average weekly consumer expenditures in the presence of calendar synchronization to average weekly consumer expenditures for the same amount of gasoline in the absence of the calendar synchronization feature. As discussed, the comparison is equivalent to a simple comparison of weekly quantity-weighted average prices on one hand and weekly unweighted average prices on the other. To control for changes in wholesale costs, I define the cost-adjusted price as the retail price minus TGP minus taxes, and present comparisons of cost-adjusted prices in the table.[8] I present weekly

[8] Cost-adjusted prices are essentially a proxy for price-cost margins.

quantity-weighted average prices in the first row of the table ("With Calendar Synchronization"), and weekly unweighted average prices in the middle row of the table ("Without Calendar Synchronization"). The main results of the article are in the last row of the table, where I compare the two ("Difference").

The results for each city in 2009-2010 and again for Perth for 2012-2013 all point to the same conclusion. Consumer expenditures are statistically significantly lower when cycles are synchronized to the calendar, compared to a counterfactual world in which the identical cycles are present but absent the calendar synchronization feature. Across all six specifications in Table 5, the average reduction in the weighted price per liter of gasoline when cycles are synchronized to the day of the week is 0.74 cents per liter, ranging from 0.34 cents per liter in Perth in 2009-2010 to 1.28 cents per liter in Adelaide in 2009-2010. The estimates are statistically significant at better than the 1% level in every case. I conclude that the intertemporal switching effect by consumers dominates any relenting-phase-shifting effect by firms in these markets, and results in more sales in periods of low prices than in periods of high prices. Consumers benefit on net from the predictability that comes with calendar synchronization.

The results agree with the conjecture made by Noel (2012) and Noel and Chu (2015) that calendar synchronization may benefit consumers by allowing them to more easily predict trough periods and purchase below the unweighted average price. It also agrees with the report of the ACCC (2007) that suggests many consumers attempt to do this. The results do not agree with the conjectures made by Noel (2007a) and Foros and Steen (2013) that calendar synchronization may harm consumers to the extent firms are able to shift relenting phases to just prior to periods of high demand.

5.3 A Falsification Exercise

Above, I attributed the difference between quantity-weighted and quantity-unweighted average prices to the presence of calendar synchronization. It is easy to confirm this interpretation by performing a falsification exercise involving the Eastern Cities – Adelaide, Brisbane, Melbourne, and Sydney – in 2012-2013. As noted above, the Eastern Cities continued to experience cycles in 2012-2013 but the calendar synchronization feature was lost by 2011.

First, I establish that cycles were present in these cities at this time but that calendar synchronization was not. Table 6 reports the derived cycle characteristics for the Eastern Cities in 2012-2013, and can be contrasted with Table 3, which reported cycle characteristics for the same cities in 2009-2010 when there was calendar synchronization. The third row of Table 6 in particular ("Cycle Period") confirms that cycles in the Eastern Cities in 2012-2013 were no longer weekly in nature. The average period of the cycle grew and now ranged from 14.1 days in Adelaide to 17.3 days in Melbourne. The standard errors were now triple or quadruple their Table 3 counterparts, reflecting the greater variation in the cycle period from one cycle to the next.

Table 7 reports results from a series of regressions of prices on day-of-the-week indicator variables and confirms that cycles in the Eastern Cities were no longer synchronized to the day of the week. In contrast to Table 4 for the earlier period with calendar synchronization, Table 7 shows no significant difference between any two prices on any two days of the week in any given city. Statistically speaking, prices were day-of-the-week invariant.

This implies that there should be no statistically significant difference between the quantity-weighted average prices and quantity-unweighted average prices in the Eastern Cities in 2012-2013, if in fact the tests are probative. To confirm this, I repeat the main comparisons of Table 5 but now for the same cities in the absence of calendar synchronization. I report the results in Table 8. The "Difference" row confirms that the falsification exercise fails as expected. It shows no difference between quantity-weighted and quantity-unweighted average prices in any city, and no change in consumer expenditures between the actual world (which in this case has no calendar synchronization) and the counterfactual world (which as always has no calendar synchronization). Table 8 confirms that, all else equal and on net, consumers in the study markets benefited from the calendar synchronization feature of the cycles.

5.4 Elasticities

As a final auxiliary analysis, I turn to an examination of gasoline price elasticities. In particular, the elasticity I am interested in is the intertemporal price elasticity of demand in response to a *temporary* change in price along the cycle. Low prices at the trough are obviously temporary and, as shown above, consumers are more likely to fill up at these times to take advantage of the temporarily low prices. This is a different type of elasticity from the more commonly estimated short and long run price elasticities of demand in response to a *permanent* change in price.[9] Estimates of intertemporal gasoline price elasticities are rare in the literature, in part because temporary "sales" on gasoline are relatively rare absent retail gasoline price cycles, and in part because high-frequency quantity data is difficult to find even when there are retail gasoline price cycles. The current application presents a unique opportunity to explore this relatively understudied effect.

I report estimates of unadjusted and adjusted intertemporal price elasticities of demand with respect to a temporary and recurring change in price in Table 9. Unadjusted elasticities are based on a simple regression of the log of volumes on the log of prices and weekly indicator variables, for all five cities during 2009-2010 and for Perth during 2012-2013. I find that these elasticity estimates are all very high and range between -6.6 and -8.7 during 2009-2010, with an estimate of -6.0 for Perth in 2012-2013. The estimates show that there is a substantial amount of intertemporal switching by consumers away from the high price days and towards the low price days of the cycle.

The unadjusted elasticity estimates are identified primarily off of the large price changes from the Wednesday price troughs to the price peaks a few days later. They contain a bias, however, because they do not control for the fact that volumes would normally be different across the days of the week even in the absence of any price differences. For example, it is well known that weekday traffic and weekday gasoline demand tends to be higher than weekend traffic and weekend gasoline demand in urban areas, all else equal. So while prices are low on Wednesdays which can increase the quantity demanded, that is also when demand would normally be high anyway. Prices are high on weekends which can decrease the quantity demanded, but that is also when demand would normally be low. The usual fix in such a case is to include day-of-the-week fixed effects in the regression to control for normal day-of-the-week differences in demand, but this is not possible in this context, since the troughs and peaks are almost perfectly correlated with the day of the week under calendar synchronization.

Fortunately, this issue is easily handled because I *do* observe what the typical underlying day-

[9] Examples include Nicol (2003), Hughes et al. (2008), Brons et al. (2008), and Lin and Prince (2013), who generally find them to be low. In the context of supermarket sales and promotions, Nevo and Hendel (2006) warn that elasticities calculated in response to a temporary reduction in price should not confused with elasticities with respect to a permanent change in price (and vice versa). The same caution applies to the interpretation of the elasticities I calculate here.

of-the-week demand pattern looks like in the study markets. As discussed above, the four Eastern Cities lost their calendar synchronization in 2011, and thereafter prices became day-of-the-week invariant, as shown in Table 7. This exposes the normal underlying day-of-the-week demand pattern in those cities. I use this pattern to adjust the relevant volumes in 2009-2010, remove the effect of normal day-of-the-week differences in demand, and estimate adjusted price elasticities free of this bias. Specifically, for each Eastern City I calculate:

$$ADJUSTED\_VOLUME_{mt} = VOLUME_{mt} - \overline{VOLUME}_{md} + \overline{VOLUME}_{m} \tag{11}$$

where $VOLUME_{mt}$ is the unadjusted volume in city m in period t during 2009-2010, $\overline{VOLUME}_{md}$ is the average volume in city m for day of the week d during 2012-2013 (with $d = \{1..7\}$), and $\overline{VOLUME}_{m}$ is the average overall volume in city m.[10] Since Perth always has calendar synchronized cycles in the data, I cannot observe its normal weekly demand pattern, and I use the average of $\overline{VOLUME}_{md}$ across the four Eastern Cities for each day d as a proxy. I use this adjusted volume in place of the raw volume when calculating adjusted elasticities.[11]

The last row in Table 9 reports estimates of adjusted elasticities. Again, the estimates are very high and contained in a narrow range from -6.3 to -6.8 across the five cities in 2009-2010 and Perth in 2012-2013. The adjusted elasticities are generally a little lower than the unadjusted elasticities, reflecting the fact that midweek volumes (when prices were relatively low) would have been a little higher than weekend volumes anyway (when prices were relatively high). More notable, however, is the fact that adjusted and unadjusted elasticities are quite similar to another, showing that the degree of bias in the unadjusted elasticities is small. In other words, the sudden rise in volumes early in the week and sudden fall in volumes late in the week is largely driven by consumers' response to cyclical pricing, rather than by normal differences in gasoline demand that differ by the day of the week. The response is strong and the intertemporal price elasticities are very high.

[10] The third term is needed to avoid negative volumes and to not exaggerate volume changes when expressed as percentages.

[11] Other factors that affect volumes but are day-of-the-week invariant are simply differenced out of the comparison and do not affect results.

6 Conclusion

In this article, I evaluated the effects of calendar synchronization on consumer expenditures. I asked whether firms were able to use calendar synchronization to move the high price days of the cycle to days when gasoline demand would normally be high anyway (the "relenting-phase-shifting effect"), or whether consumers were able to use the more predictable pattern of the cycle to effectively shift purchases to low price days ("the intertemporal switching effect"). I found that the intertemporal switching effect by consumers dominates the potential relenting-phase shifting effect by firms, and consumer expenditures fall with calendar synchronization, holding all else equal. I discussed why it remains a reasonable course of action for certain firms to try to move to a calendar synchronized cycle, and others to follow this move, even if the net effect is that consumers gain overall.

The main results are based on comparing weekly average quantity-weighted prices in the presence of calendar synchronization to weekly average quantity-weighted prices for the same cycles without calendar synchronization. Since prices are day-of-the-week invariant in the absence of calendar synchronization, quantity weights are irrelevant in the counterfactual world and the comparison boils down to a simple comparison of weekly quantity-weighted average prices and weekly unweighted average prices. It is a simple methodology that can be implemented by various competition authorities and others to evaluate the volume shifting effects of calendar synchronization in their own jurisdictions.

I also calculate intertemporal price elasticities in response to a recurring and temporary change in the price of gasoline. These types of elasticities are rarely encountered in the literature because temporary "sales" on gasoline are uncommon in the absence of retail gasoline price cycles, and high-frequency volume data is difficult to come by even with retail gasoline price cycles. I find elasticities to be very high, confirming that consumers are responsive to the cycle and in large numbers shift their purchases to known low price days. While several other articles have conjectured that consumers could do this (Noel (2012), Noel and Chu (2015)), the results of this article show that consumers actually *do* do this. A second implication of the results is that simple quantity-unweighted average prices – as is often calculated in the literature on retail gasoline price cycles and used in various price comparisons – tend to overstate the prices that consumers really pay when there are retail gasoline price cycles.

Finally, this article emphasizes the importance of using gasoline volume data whenever possible. The lack of high-frequency gasoline volume data is an ongoing challenge in the literature, and this article represents a rare exception in this regard. In emphasizing the value of volume data, this article highlights an important direction forward for future gasoline research. Understanding how gasoline volumes and gasoline prices interact in both the short run and the long run has important implications for policy and welfare, and for economists seeking to better understand the competitive forces at play in the retail gasoline industry.

7 References

ACCC: Australian Competition and Consumer Commission (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. (2009). "Retail Gasoline Price Cycles: Evidence from Guelph, Ontario Using Bi-Hourly, Station-Specific Retail Price Data", *Energy Journal* 30, 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.

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.

Brons, M., P. Nikkamp, E. Pels and P. Rietveld (2008), "A Meta-Analysis of the Price Elasticity of Gasoline Demand. A SUR Approach", *Energy Economics*, 30(5), 2105-2122.

Byrne, D. and N. de Roos (2017). "Learning to Coordinate: A Study in Retail Gasoline", University of Melbourne and University of Sydney working paper.

Byrne, D.P. and N. de Roos (2015). "Consumer Search in Retail Gasoline Markets", *Journal of*

Industrial Economics, 183-193.

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.

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

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

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

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

Haucap, J., U. Heimeshoff, and M. Siekmann (2016). "Selling Gasoline as a By-Product: The Impact of Market Structure on Local Prices", DICE Working Paper 240, University of Dusseldorf, 1-33.

Hughes, J.E., C.R. Knittel, and D. Sperling (2008), "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand", *Energy Journal*, 29 (1), 113-134.

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), 581-606.

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

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.

Lin, C.-Y.C. and L. Prince (2013), "Gasoline Price Volatility and the Elasticity for Gasoline", *Energy Economics* 38, 111-117.

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

Nevo, A. and I. Hendel (2006). "Measuring the Implications of Sales and Consumer Inventory Behavior", *Econometrica* 74(6), 1637-1673.

Nicol, C.J. (2003), "Elasticities of Demand for Gasoline in Canada and the United States", *Energy Economics*, 25(2), 201-214.

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, 324-334.

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

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

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

Noel, M.D. (2015). "Do Edgeworth Price Cycles Lead to Higher or Lower Prices?", *International Journal of Industrial Organization* 42, 81-93.

Noel, M.D. and L. Chu (2015). "Forecasting Gasoline Prices in the Presence of Edgeworth Price Cycles", *Energy Economics* 51, 204-214.

Siekmann (2017). "Characteristics, Causes, and Price Effects: Empirical Evidence of Intraday Edgeworth Cycles", DICE working paper 252, University of Dusseldorf, 1-35.

Valadkhani, A. (2013). "Seasonal Patterns of Daily Prices of Unleaded Petrol Across Australia", *Energy Policy* 56, 720-731.

Wang, Z. (2009a). "Station Level Gasoline Demand in an Australian Market with Regular Price Cycles", *Australian Journal of Agricultural and Resource Economics* 53, 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, 987-1030.

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. Example of Retail Gasoline Price Cycles, Perth in 2012-2013

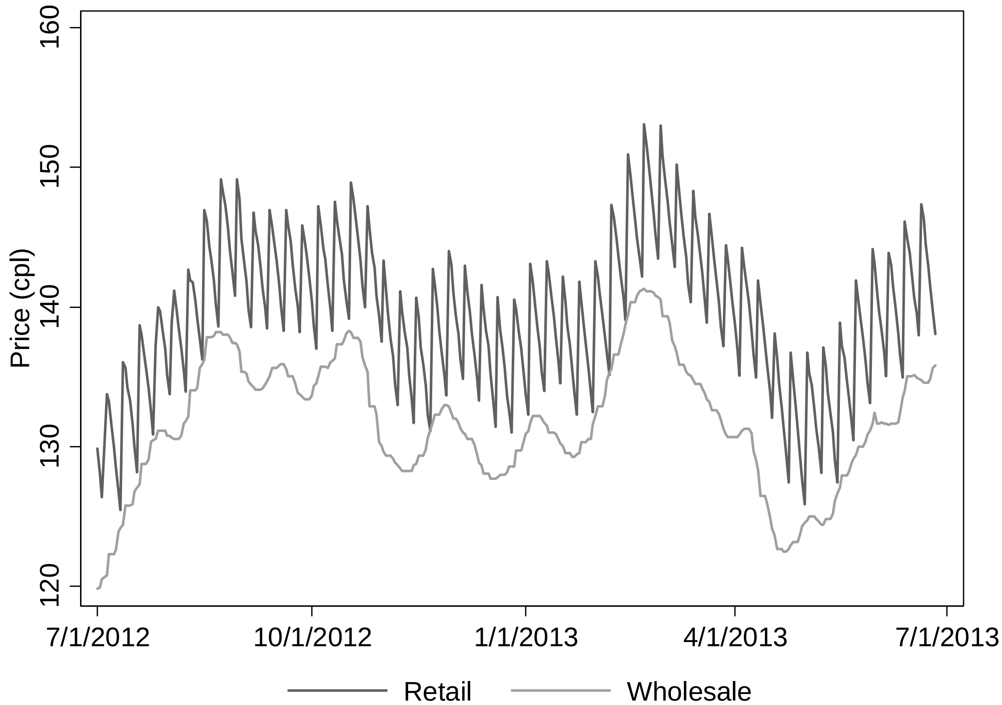

Figure 2. Volumes by Day of the Week with Calendar Synchronization

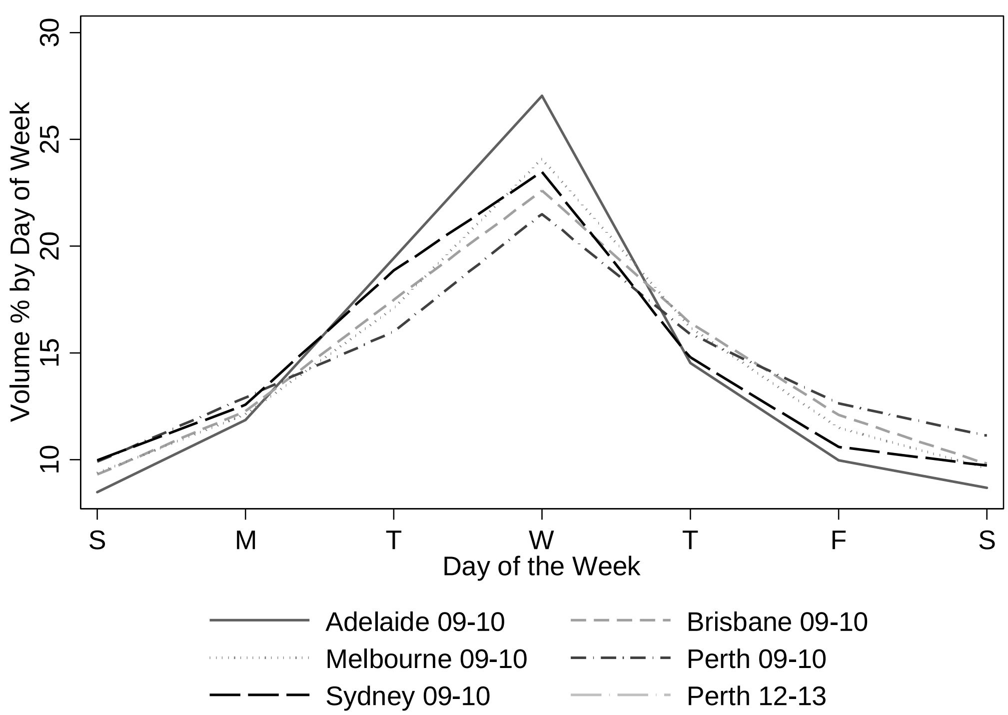

Table 1. Summary Statistics

	Num. Obs.	Mean	Std. Dev.	Minimum	Maximum
Retail Price	3650	135.31	7.12	115.6	155.3
Wholesale Price	3650	128.10	5.53	114.5	141.3
Cost-Adjusted Price	3650	7.21	5.06	(4.9)	19.0
Volume	3650	14.29	3.74	8.5	27.0

Table 2. Core Cycle Parameters

	(1) Adelaide 2009-2010	(2) Brisbane 2009-2010	(3) Melbourne 2009-2010	(4) Perth 2009-2010	(5) Sydney 2009-2010	(6) Perth 2012-2013
Within-Phase Price Changes						
α^R = E($\Delta RETAIL_{mt}$ | I_{mt} = "R")	7.231**	4.975**	5.898**	2.860**	6.216**	8.925**
(average price change - relenting)	(0.516)	(0.253)	(0.321)	(0.210)	(0.335)	(0.278)
α^U = E($\Delta RETAIL_{mt}$ | I_{mt} = "U")	-2.671**	-1.890**	-2.218**	-1.313**	-2.174**	-1.507**
(average price change - undercutting)	(0.098)	(0.052)	(0.056)	(0.064)	(0.079)	(0.019)
Phase Transition Probabilities						
λ^{RR} (R to R)	0.470**	0.490**	0.495**	0.591**	0.479**	0.036
	(0.050)	(0.050)	(0.050)	(0.046)	(0.051)	(0.025)
λ^{RU} (R to U)	0.530**	0.510**	0.505**	0.409**	0.521**	0.964**
	(0.050)	(0.050)	(0.050)	(0.046)	(0.051)	(0.025)
λ^{UR} (U to R)	0.195**	0.193**	0.189**	0.183**	0.181**	0.167**
	(0.024)	(0.024)	(0.024)	(0.024)	(0.023)	(0.021)
λ^{UU} (U to U)	0.805**	0.807**	0.811**	0.817**	0.819**	0.833**
	(0.024)	(0.024)	(0.024)	(0.024)	(0.023)	(0.021)
Num. Obs.	365	365	365	365	365	365

** Significant at the 5% level. * Significant at the 10% level. Robust standard errors in parentheses.

Table 3. Derived Cycle Characteristics

	(1) Adelaide 2009-2010	(2) Brisbane 2009-2010	(3) Melbourne 2009-2010	(4) Perth 2009-2010	(5) Sydney 2009-2010	(6) Perth 2012-2013
Relenting Phase Duration	1.887**	1.962**	1.980**	2.447**	1.920**	1.038**
	(0.178)	(0.191)	(0.195)	(0.275)	(0.188)	(0.027)
Undercutting Phase Duration	5.115**	5.176**	5.300**	5.457**	5.510**	5.981**
	(0.637)	(0.652)	(0.676)	(0.728)	(0.713)	(0.758)
Cycle Period	7.002**	7.138**	7.280**	7.903**	7.430**	7.019**
	(0.661)	(0.679)	(0.703)	(0.777)	(0.737)	(0.758)
Cycle Amplitude	13.644**	9.759**	11.680**	6.998**	11.936**	9.262**
	(1.399)	(1.045)	(1.316)	(0.933)	(1.263)	(0.245)
Cycle Asymmetry	2.711**	2.639**	2.676**	2.230**	2.870**	5.763**
	(0.424)	(0.420)	(0.432)	(0.390)	(0.467)	(0.746)
Num. Obs.	365	365	365	365	365	365

** Significant at the 5% level. * Significant at the 10% level. Durations and cycle periods in days, amplitudes in cents per liter, asymmetries are unit free. Robust standard errors in parentheses.

Table 4. Prices by Day of the Week

Dep. Var.: Price	(1) Adelaide 2009-2010	(2) Brisbane 2009-2010	(3) Melbourne 2009-2010	(4) Perth 2009-2010	(5) Sydney 2009-2010	(6) Perth 2012-2013
Monday	-2.600**	-1.988**	-2.332**	-1.713**	-2.318**	-1.703*
	(1.085)	(0.865)	(0.911)	(0.846)	(0.918)	(0.868)
Tuesday	-5.315**	-4.242**	-4.848**	-3.451**	-4.582**	-3.360**
	(1.085)	(0.865)	(0.911)	(0.846)	(0.918)	(0.868)
Wednesday	-7.048**	-6.006**	-6.877**	-3.430**	-6.150**	-4.922**
	(1.080)	(0.861)	(0.907)	(0.842)	(0.914)	(0.868)
Thursday	2.108*	-3.103**	-3.778**	-0.665	-1.625*	4.247**
	(1.085)	(0.865)	(0.911)	(0.846)	(0.918)	(0.868)
Friday	4.286**	1.014	1.146	2.453**	2.871**	3.070**
	(1.085)	(0.865)	(0.911)	(0.846)	(0.918)	(0.868)
Saturday	2.331**	1.359	1.555*	1.450*	2.124**	1.601*
	(1.085)	(0.865)	(0.911)	(0.846)	(0.918)	(0.868)
Constant	132.500**	135.573**	135.395**	130.105**	132.298**	139.612**
	(0.767)	(0.611)	(0.644)	(0.598)	(0.649)	(0.611)
Num. Obs.	365	365	365	365	365	365

** Significant at the 5% level. * Significant at the 10% level. Robust standard errors in parentheses.

Table 5. Comparison of Cost-Adjusted Prices With and Without Calendar Synchronization

Dep. Var.: Cost-Adjusted Prices	(1) Adelaide 2009-2010	(2) Brisbane 2009-2010	(3) Melbourne 2009-2010	(4) Perth 2009-2010	(5) Sydney 2009-2010	(6) Perth 2012-2013
With Calendar Synchronization	4.783**	8.414**	7.539**	3.657**	5.380**	6.484**
(Quantity-weighted)	(0.214)	(0.148)	(0.175)	(0.158)	(0.182)	(0.150)
Without Calendar Synchronization	6.065**	9.081**	8.380**	4.001**	6.232**	6.963**
(Unweighted)	(0.267)	(0.188)	(0.224)	(0.175)	(0.231)	(0.161)
Difference	-1.282**	-0.667**	-0.841**	-0.344**	-0.852**	-0.479**
	(0.111)	(0.136)	(0.134)	(0.053)	(0.100)	(0.072)
Num. Obs.	365	365	365	365	365	365

** Significant at the 5% level. * Significant at the 10% level. Robust standard errors in parentheses.

Table 6. Falsification Exercise: Derived Cycle Characteristics in Cities Without Calendar Synchronization

	(7)	(8)	(9)	(10)
	Adelaide	Brisbane	Melbourne	Sydney
	2012-2013	2012-2013	2012-2013	2012-2013
Relenting Phase Duration	2.462**	4.125**	3.545**	4.043**
	(0.372)	(0.734)	(0.641)	(0.732)
Undercutting Phase Duration	11.615**	11.609**	13.714**	11.870**
	(2.181)	(2.317)	(2.885)	(2.372)
Cycle Period	14.077**	15.734**	17.260**	15.913**
	(2.211)	(2.428)	(2.954)	(2.480)
Cycle Amplitude	15.813**	10.844**	12.940**	12.646**
	(2.856)	(2.186)	(2.801)	(2.699)
Cycle Asymmetry	4.719**	2.814**	3.868**	2.935**
	(1.139)	(0.754)	(1.075)	(0.793)
Num. Obs.	365	365	365	365

** Significant at the 5% level. * Significant at the 10% level. Durations and cycle periods in days, amplitudes in cents per liter, asymmetries are unit free. Robust standard errors in parentheses.

Table 7. Falsification Exercise: Prices by Day of the Week in Cities Without Calendar Synchronization

Dep. Var.: Price	(7)	(8)	(9)	(10)
	Adelaide	Brisbane	Melbourne	Sydney
	2012-2013	2012-2013	2012-2013	2012-2013
Monday	-0.373	-0.117	-0.543	-0.728
	(1.368)	(1.188)	(1.331)	(1.253)
Tuesday	-0.356	-0.558	-0.925	-1.312
	(1.368)	(1.188)	(1.331)	(1.253)
Wednesday	0.024	-0.592	-1.141	-1.200
	(1.368)	(1.188)	(1.331)	(1.253)
Thursday	0.445	-0.330	-0.943	-0.439
	(1.368)	(1.188)	(1.331)	(1.253)
Friday	-0.168	-0.266	-0.771	-0.070
	(1.368)	(1.188)	(1.331)	(1.253)
Saturday	-0.594	-0.269	-0.319	-0.052
	(1.368)	(1.188)	(1.331)	(1.253)
Constant	137.996**	142.035**	138.435**	138.048**
	(0.962)	(0.836)	(0.937)	(0.882)
Num. Obs.	365	365	365	365

** Significant at the 5% level. * Significant at the 10% level. Robust standard errors in parentheses.

Table 8. Falsification Exercise: Comparison of Cost-Adjusted Prices Without Calendar Synchronization in 2012-2013

Dep. Var.: Cost-Adjusted Prices	(7)	(8)	(9)	(10)
	Adelaide	Brisbane	Melbourne	Sydney
	2012-2013	2012-2013	2012-2013	2012-2013
Quantity-weighted	5.234**	8.892**	5.460**	5.178**
	(0.291)	(0.209)	(0.247)	(0.233)
Unweighted	5.229**	8.923**	5.499**	5.198**
	(0.288)	(0.202)	(0.249)	(0.235)
Difference	0.005	-0.030	-0.039	-0.020
	(0.049)	(0.066)	(0.039)	(0.034)
Num. Obs.	365	365	365	365

** Significant at the 5% level. * Significant at the 10% level. Robust standard errors in parentheses.

Table 9. Adjusted and Unadjusted Intertemporal Price Elasticities in Response to a Temporary Change in Price

	(1)	(2)	(3)	(4)	(5)	(6)
	Adelaide	Brisbane	Melbourne	Perth	Sydney	Perth
	2009-2010	2009-2010	2009-2010	2009-2010	2009-2010	2012-2013
Unadjusted Elasticities	-6.643**	-8.655**	-7.719**	-7.646**	-7.195**	-6.020**
	(0.451)	(0.453)	(0.430)	(0.524)	(0.363)	(0.567)
Adjusted Elasticities	-6.426**	-6.846**	-6.312**	-6.836**	-6.821**	-6.656**
	(0.358)	(0.388)	(0.375)	(0.358)	(0.342)	(0.372)
Weekly Indicator Variables	Y	Y	Y	Y	Y	Y
Num. Obs.	365	365	365	365	365	365

** Significant at the 5% level. * Significant at the 10% level. Robust standard errors in parentheses.